T-CAIREM's Healthcare AI Governance Framework

Operationalizing AI Governance in Learning Health Systems

Abhishek Moturu, Gemma Postill, Laura Rosella, Muhammad Mamdani

Temerty Centre for Artificial Intelligence Research and Education in Medicine
Toronto Academic Health Science Network AI in Healthcare Community of Practice
University of Toronto

1 October 2026

Temerty Centre for AI Research and Education in Medicine
UNIVERSITY OF TORONTO

Toronto Academic Health Science Network

Contents

1 Summary 3

2 How to use this framework 3

3 Purpose of hospital AI governance 5

4 Evidence and framework development 6

5 Required governance functions 7

6 AI inventory and use boundaries 9

7 Accountability and decision authority 11

8 Sourcing, procurement, and change control 16

9 Risk tiering and evidence for approval 22

10 Monitoring and lifecycle decisions 27

11 Literacy, transparency, and participation 39

12 Governance implementation and evaluation 45

References 50

Appendix A Worked examples 60

Appendix B Minimum governance model checklist 88

Appendix C Comprehensive governance model checklist 96

Appendix D Evidence requirements by risk tier 102

Appendix E Lifecycle implementation checklist 109

Appendix F Checklists for system capabilities 125

Appendix G Procurement and external data flows 133

Appendix H AI incident learning record 139

Appendix I Shared fields for governance records 141

Appendix J Evidence sources and local adaptation 153

Appendix K Glossary 170

Appendix L Hospital AI Governance Intake and Lifecycle Record 178

1 Summary

Artificial intelligence is playing a growing role in how hospitals deliver care and manage their operations. Its value depends on whether the benefits justify the risks and whether hospitals can oversee its use as circumstances change. Governance assigns responsibility for making these judgments and acting on them. Evaluation must therefore examine how AI affects care and work over time, as well as how well the model performs. [3][4] This approach is consistent with a learning health system, in which evidence guides practice and experience from routine care informs further evaluation and improvement. [1][2] This framework explains how hospitals can organize that work.

The framework covers AI used in clinical care and hospital operations, including decision support, imaging, and pathology. It also covers ambient documentation, which captures clinical conversations, and generative AI, which produces text, images, or other content. Agentic AI is included where a system plans or carries out tasks using software tools. An AI model is the algorithm that produces an output. The AI system includes that model, its data and interfaces, and the people, oversight arrangements, and workflow through which its outputs affect care or work. Assessment covers how people experience the service and its consequences, including those not captured by model performance.

Any hospital or hospital group can use the framework, with adaptations for its jurisdiction and decision-making arrangements. Its development draws on the literature and guidance discussed in **Appendix J**, including research on how people use technology under actual working conditions. Consultation with representatives across all 15 Toronto Academic Health Science Network (TAHSN) member organizations and with other health-AI leaders informed the procedures and examples.

Hospitals can organize oversight through existing staff and committees, dedicated AI governance teams, or shared expertise across hospitals. The minimum and comprehensive models describe these differences in staffing and resources. Both require the same protections for a use with the same risks. A hospital with a small team may oversee a higher-risk service if it can provide the required evaluation, specialist review, monitoring, and response to problems. Everyone using or overseeing the service needs the knowledge and skills required for their role, referred to here as AI literacy. Conversely, a larger team need not repeat assessments that add no relevant evidence for a familiar lower-risk use.

The framework's practical contribution is to connect hospital decisions with a record that staff can use from the first proposal until the AI service is retired. It links each approval condition to a responsible person, evidence that the condition has been met, and the action required if it has not. The checks cover discriminative, generative, and agentic AI. Hospitals can share assessments and specialist expertise, while each remains responsible for decisions about its own use of the system. **Appendix J** explains how this approach draws on earlier work and what it adds.

2 How to use this framework

A hospital can use this framework independently or with other hospitals. Each hospital documents how the shared framework applies to its own organization in a hospital-specific appendix. This states which decisions are made by a shared governance body, if there is one, and which remain the hospital's responsibility. Leaders first decide how existing committees, dedicated staff, or a shared service will provide oversight. The choice affects how work is organized, but does not reduce the evidence or safety requirements for an individual use.

Review starts with Tier 0, an initial screen of whether the next proposed activity may proceed. A prohibited or unacceptable use is rejected; a proposal with missing essential information or an unresolved prerequisite is held for clarification or correction. Once there is enough information to assess the use, reviewers assign a provisional

Tier 1–4 according to the most serious harm that could plausibly result. The tier determines how much evidence and oversight are needed. The intended use and enabled features determine which specialist reviews, staff competencies, and safeguards apply. **Section 9** explains how to make these judgments and record who may authorize the next step.

The records serve different purposes. The intake form describes the proposed use and why it is needed. The AI registry is the hospital's inventory of AI uses and their current authorized status. The AI dossier is the evidence and decision file for one use; it can link to existing documents rather than duplicate them. Governance approval states which activities are permitted and under what conditions. Before live use, production confirmation checks that the actual version and setup meet that approval and that every requirement due before use has been met. Monitoring findings and later decisions are added to the same record, including changes, incidents, and retirement. **Appendices D–I** give the detailed requirements, and **Appendix L** provides the adaptable Hospital AI Governance Intake and Lifecycle Record.

Table 1 shows the seven decision stages that take a proposal from initial review to operation and retirement. **Figure 5** groups the work needed to support those stages into six recurring tasks; it is a guide to the work, not another approval sequence. The service's current status is recorded separately from the decision being considered. For example, a service may remain in routine operation while reviewers assess a proposed expansion. An incident or change may require an earlier decision to be reconsidered without changing every part of the service's approval.

For a new proposal, the requester starts with **Sections 1–3** of the record in **Appendix L**. These sections describe the problem, current workflow, people affected, proposed information flow, alternatives, and expected benefit. A complete technical file is not required at first contact. The governance lead helps define the use, assigns specialist questions, and coordinates the review. A dossier coordinator keeps the evidence and decisions linked as the review progresses; this administrative role does not carry approval authority. **Appendix I** defines the fields used across the records, and **Appendix K** explains unfamiliar terms.

A checklist item is Met when the reviewer has accepted evidence that the requirement is satisfied for the decision under review. Receiving a document or completing a test is not enough if its findings are unacceptable. An Open item is marked Pending, Failed, Unknown, or Expired and records who must act, what they must do, the due date, and the decision or activity that must wait. A requirement due later may remain Open—Pending while earlier work proceeds, provided that work does not depend on it. N/A requires a recorded reason accepted by the reviewer. It cannot waive a mandatory safeguard.

The readiness assessment summarizes these findings for the requested activity. Ready means all requirements due for that decision have been met. Ready with required actions permits only the work explicitly authorized to proceed while other items remain open; every open item needs an owner and due date. Not ready means the requested activity cannot proceed. The authorized decision-maker uses this assessment to decide what may happen next. A completed checklist or a Ready rating informs that decision; neither grants governance approval or permission for live release.

The appendices support different tasks and need not all be completed at intake. **Appendix A** provides worked examples. **Appendices B and C** help hospitals establish their governance arrangements; **Appendices D–F** provide checks for risk tiers, stages of use, and system capabilities. **Appendices G–I** support procurement, incident review, and consistent records. **Appendix J** explains the evidence and local adaptations, **Appendix K** is the glossary, and **Appendix L** is the adaptable record for an individual AI use. When several records rely on the same evidence, such as an evaluation report, they can refer to one stored copy.

3 Purpose of hospital AI governance

Hospital AI governance establishes who may decide how AI is used and what evidence they need. It starts with the problem the hospital wants to address and continues through procurement, evaluation, introduction, changes, and eventual retirement. During use, the hospital must monitor safety, effectiveness, equity, and responsible practice. If the grounds for approval no longer hold, it must be able to investigate and, where necessary, restrict or end the service.

AI may enter a hospital through a regulated device, an internally developed model, a vendor service, an embedded software feature, or a general-purpose generative system. It does not necessarily begin as a formal AI proposal. [11][3][4] Software updates, departmental purchases, research activity, and staff accounts can introduce uses that a review of named products would miss. Staff therefore need accessible advice and a list of approved AI services showing where and by whom each may be used, and which data and tasks are permitted. These resources also provide a route for disclosing unregistered use before harm occurs. The US National Institute of Standards and Technology (NIST) Generative AI Profile and work on the 'AI iceberg' inform the framework's approach to maintaining that inventory, setting acceptable-use rules and limits, and overseeing suppliers. [12][13]

The depth of review should reflect the harms that could plausibly follow from the proposed use. These can arise in clinical care and communication or through access, scheduling, staffing, employment, and resource allocation. [6][3][14] Reviewers must judge what evidence is sufficient for the setting, whether human oversight can work under ordinary conditions, and how privacy, cybersecurity, safety, equity, and continuity will be protected. [15]

Those judgments may need to be revisited even when a product retains its name. Approval rests on assumptions about the data pipeline—the steps that collect, prepare, and move data—and about prompts, decision thresholds, interfaces, and safeguards. Changes to these, or to staffing, workflows, and supplier versions, can affect the basis of approval. [16][17][18] The responsible owners must document such changes and seek review of any assumptions that may no longer hold.

Responsibility for an AI service is shared between its clinical and operational owners and the hospital functions authorized to oversee data, technology, safety, law, procurement, research, workforce matters, and institutional risk. [19][20][14] Their required findings must be retained in the decision record and considered together, including findings about how AI redistributes work. Time saved by one team may require another to check or correct outputs, explain results, or handle exceptions. Staffing changes and fewer opportunities to practise may also erode skills needed for safe use. Reviewers must examine these effects. The staff, patients, families, and communities affected need a way to contribute evidence and influence the decision.

When AI materially changes an interaction, patients and clinicians need understandable information about the system's capabilities and limits, how information is handled, who remains responsible, and how to raise a concern or report an error. [6][21][22] This information should accompany the technical documentation and support people in questioning use and seeking correction. The hospital must provide a way to contest a consequential outcome and a practical alternative where feasible or required, as described in **Section 11.4**. These measures make the hospital's responsibilities clearer to the people affected by its decisions and help sustain public and professional trust.

National approaches offer different forms of support. Canada's 2026 National Artificial Intelligence Strategy, AI for All, makes health care the first focus of its AI Missions Program and emphasizes trust, literacy, safeguards, and Canadian-controlled data and infrastructure. [23] Singapore's AIHGle 2.0 sets out responsibilities for AI developers, healthcare organizations, and users. [8] In the United States, the voluntary National Institute of Standards and Technology (NIST) AI Risk Management Framework addresses risks

throughout development and use, and the Coalition for Health AI (CHAI) provides guidance for responsible AI. [42][9] The UK National Institute for Health and Care Excellence (NICE) sets out evidence expectations for digital health technologies. [10] Hospitals can use this guidance to organize oversight, but still need evidence that each proposed use is lawful, safe, effective, equitable, and workable locally. This includes arrangements for monitoring, incidents, and changes.

4 Evidence and framework development

Existing guidance gives hospitals principles and assessment methods, but often leaves them to work out how committees, specialist teams, clinical services, procurement, and quality programs should apply them. [5][6] Wang, Freeman, and Magrabi's 2026 scoping review found that only 10 of 77 frameworks combined principles, an assessment method, coverage throughout an AI system's life, and a mechanism for oversight. Hospitals still had to determine who should review a proposal and how findings should inform purchasing, the start of live use, and later changes. [5]

A separate systematic review developed the Healthcare AI Governance Readiness Assessment (HAIRA), which examines whether the hospital has the capabilities needed for governance and how those capabilities depend on one another. An essential missing capability limits its assessed readiness. [7] For example, evidence that an AI system performs adequately does not establish that the hospital can operate it safely. The hospital also needs records, incident response, a safe alternative when AI is unavailable or unsuitable (fallback), and procedures to assess and authorize changes before they are introduced (change control). The decision record should show whether these arrangements support the proposed activity and identify unresolved issues that prevent it from proceeding.

Hospitals need enough staff and expertise to review their AI uses and oversee them after approval. Excessive delays can encourage informal workarounds, while unclear responsibility can leave problems unaddressed. Staffing and committee arrangements may vary with the number and complexity of uses, provided every required function has someone accountable for it. Shared assessment methods, procurement terms, education, monitoring, and incident learning can reduce duplicated work. Each hospital must still determine whether the shared expertise and resources meet its needs and preserve its authority over local decisions.

Feedback from TAHSN survey respondents informed the framework's treatment of staffing and resources. Respondents described AI steering committees or equivalent arrangements overseeing several kinds of AI. They reported different levels of governance maturity, limited capacity, competing priorities, uneven AI literacy, and rapid technical change. Their accounts informed the proposed shared records, education, reviewer support, and procurement requirements. This survey feedback concerns the responding organizations; the broader consultation across 15 TAHSN institutions is described in **Appendix J2**. Hospitals can use shared resources and add dedicated staff as the number or complexity of their AI services grows. Throughout that growth, every required function must remain covered, with clear responsibility and sufficient expertise for consequential uses.

Hospitals should use experience with AI to inform later approvals, purchasing, and education. The record should cover benefits as well as failures, near misses, workflow changes, inequitable effects, and reasons for retirement. Each finding should be linked to the version and setup in use and the decision that permitted it. Changes to governance procedures should identify the practical evidence that prompted them.

Evidence about organizational AI governance remains uneven. Reviews and standards identify common responsibilities, and hospital reports describe ways to organize them. Few studies assess governance as a whole or its effects on patients, staff, equity, and hospital outcomes. **Appendix J** separates empirical findings from practice guidance and identifies recommendations made by this framework. Hospitals should evaluate those recommendations during implementation to learn how well they work in different settings.

Public investment can help hospitals obtain research infrastructure, data, and expertise; **Appendix J8** describes examples from several countries. Implementation budgets must also cover integration, secure data access, independent evaluation, staff training, patient involvement, monitoring, and an affordable way to stop or replace the service. Participation in a funded initiative does not establish clinical benefit. The hospital must assess the proposed use in its own population and workflow and determine whether it can sustain the service.

5 Required governance functions

The hospital must be able to show what AI is being used or considered, who is responsible for it, and why its use was approved. Before release, each use needs named accountability for clinical or end-user decisions, technical operation, and day-to-day workflow, alongside the functions accountable for privacy, cybersecurity, and patient safety. The approval record defines what is permitted and links it to the supporting evidence. It also states the value case: why the expected benefits justify the costs, workload, and risks. Reviewers use this record to check whether the service in operation still matches what was approved.

These responsibilities continue after the service starts. Monitoring must identify when a result requires action, and staff must know how to report and respond to incidents. Contracts and working arrangements must support ongoing oversight, whether the hospital develops the AI, buys it, or works with a partner. Planning must cover updates and eventual retirement as well as the initial introduction. Experience with one service should inform decisions about other AI uses.

The hospital must also prepare people for their roles, explain relevant AI uses, and provide ways to question decisions or seek correction. Lessons must be shared across sites and used to keep local practice and the governance framework current. The hospital maintains an approved set of current templates, standards, procurement terms, and educational resources so staff can find the instructions that apply. These responsibilities remain necessary wherever the committee sits within the hospital.

5.1 Minimum governance model

A hospital needs an AI Steering Committee or a formally designated equivalent. Its mandate must cover receiving requests, conducting Tier 0 screening and tiering, coordinating specialist review, recording decisions, and referring matters beyond its authority. An existing governance body can perform these duties if it has the mandate and capacity; a new department is not necessary. The board or equivalent governing body sets the hospital's risk appetite: the kinds of risk it is willing to consider in pursuing benefits. It also sets the limits beyond which a decision must be referred to a higher authority. Within its formally assigned authority, the committee decides on individual uses and records the evidence and conditions supporting each decision.

Every approved use requires the three owners defined in **Section 7**. The clinical or end-user owner is accountable for appropriate use and outcomes. The internal technical owner is accountable for the technical service operated locally, and the operational owner for workflow, support, and sustainability. Together, the intake, registry entry, and dossier record what is authorized, why it was approved, and any outstanding conditions. They also show whether the version, setup, and working conditions in actual use still match the approval and how monitoring, changes, and retirement will be managed. Other hospital functions contribute under their existing authority, and their required findings remain part of the record.

Existing staff can carry out the minimum model's duties part-time if each duty has an accountable owner and the records preserve continuity when staff change. The arrangements must cover movement from research into routine use, supplier performance, user preparation, and incident response. Shared or regional teams may provide reviewers, evaluation methods, and education. The hospital must still decide whether the service suits

the affected population, workflow, resources, and data environment, and whether the remaining risk is within the limits it is authorized to accept.

Governance approval records the permitted use, supporting evidence, conditions, and accepted residual risk—the risk that remains after safeguards. The decision-maker must explicitly accept that risk within their assigned authority. Before live use, production confirmation checks that the actual version and setup meet the approval, including data routes and connections to other systems, prepared users, monitoring, support, and fallback. The same check is required for the affected part of an approved expansion, a material change to the live system, or a restart. **Appendix B** gives the minimum governance checklist.

5.2 Comprehensive governance model

The comprehensive model retains all minimum requirements and provides dedicated staff and ongoing resources for a larger or more complex set of AI uses. The AI Steering Committee operates under a written charter defining its authority, meeting schedule, quorum (the attendance required for a valid decision), management of conflicts, decision records, and reporting to executive leadership. It reviews individual uses, incidents across hospital AI services, supplier dependencies, recurring funding, and available capacity. It considers how these decisions affect other services and institutional priorities.

The committee's regular membership should include clinical and operational leaders and representatives responsible for technology, privacy, cybersecurity, quality, legal obligations, and procurement. Frontline staff need a reliable way to contribute, either through membership or an established representative. Patients and families, rights and equity specialists, and representatives of education, research, finance, communications, and organizational risk participate when a use requires their knowledge or falls within their mandate. The committee records their conclusions and preserves each reviewer's authority. It refers unresolved risks to the person authorized to decide them.

For each use, the three owners assemble a Project Working Group with the relevant specialists to prepare, implement, and monitor the service. The governance lead coordinates intake, evidence requirements, review assignments, and follow-up of conditions. The dossier coordinator keeps the evidence file, registry entry, and decision history linked and current. Both roles need a named person to cover absences or transitions. A delayed review or staff change must not leave a live service without oversight or an outstanding condition without follow-up.

Hospitals must specify who takes responsibility when work moves between procurement, research, implementation, monitoring, and change review. Everyone needs preparation for their assigned role, and findings from actual use should inform training and practice. Dedicated educators, regular review across sites, and systematic analysis can help sustain this work as the number of AI uses grows.

A hospital-wide dashboard helps the committee identify services needing attention. It should show each service's approved status and highlight overdue reviews, incidents, material changes, supplier dependencies, and use outside approved limits. Renewal dates for contracts and recurring funding should be included so approaching expiry of evidence, rights, or resources is visible. The registry identifies each use and its owners; the linked dossier contains the evidence needed to investigate a concern or make a decision.

Email and spreadsheets may be adequate for a small number of AI uses. The hospital must still specify when the volume of requests, changes, incidents, or overdue actions will require dedicated staff or software to assign and track governance work. Once that threshold is reached, leadership must fund more capacity, slow new work, or reduce the number of different AI setups it supports. Software can assign reviews and keep an audit trail of what was done, by whom, and when. Qualified reviewers remain responsible for Tier 0 decisions, final consequence

tiers, and judgments about residual risk. The software itself requires an owner, arrangements for continuity, and a way to export records if it is replaced. **Appendix C** provides the comprehensive model checklist.

6 AI inventory and use boundaries

Approval applies to a particular use of an AI system. A platform approved for administrative drafting may need a different assessment before it provides clinical advice, even if the model has not changed. Reviewers need to know who will use the output and what decisions or actions may follow. The description of the use must identify its purpose, users, patient or other affected population, setting, and permitted features. These details define its scope: the limits within which approval applies.

Configuration means the version of the AI model or software product and how the hospital sets it up for the intended use. It includes settings and prompts (instructions given to the model), connections to data and other software, user interfaces, and the surrounding workflow and safeguards. The hospital should assess the features and workflows it plans to use. A licence covering the whole hospital, regulatory authorization, or favourable evaluation can inform the assessment, but does not establish that every feature or use is appropriate. For example, approval to draft clinical notes does not automatically cover using the same product to recommend treatment.

The appropriate level of AI involvement depends on the task and setting. Each task needs a separate description; one service may use several of the following arrangements. For clinical tasks, the responsible person is an appropriately qualified clinician. For other hospital work, it is an authorized staff member.

1. Abstain (human-only): A person performs the task without AI assistance.

2. Suggest: AI provides a recommendation. The responsible person considers it and decides what to do.

3. Draft: AI prepares work that the responsible person reviews and approves before it is used or communicated.

4. Act with direct oversight: AI carries out the task while the responsible person supervises and can intervene in time to prevent or limit harm.

5. Carry out defined tasks independently: AI acts within approved limits without prior approval of each action. Performance is monitored, with specified conditions for human review or escalation.

Here, abstain means keeping a task human-only; it does not mean a model withholding an output. These arrangements describe the AI's role, not its risk tier. Even a recommendation can have serious consequences. Assessment under **Section 9** addresses the consequences and effectiveness of oversight; the arrangement for each task is recorded in **Table L17**.

The description of each proposed use explains how the AI output reaches a patient, staff member, learner, or operational team and what they may do with it. It records the purpose, users, affected population, setting, and software or model version. It explains what shapes the output, including data sources, prompts, retrieved information, decision thresholds, and rules, and traces the output through the workflow to the decisions or actions it can affect. The account must also identify supplier duties, human authority, monitoring, and fallback. Evidence from development or regulatory review can inform assessment, but reviewers must check whether it applies to the proposed hospital, population, data connections, setup, and workflow.

A single platform may be used to draft non-sensitive internal text, summarize charts, prepare patient messages, rank applicants, or arrange follow-up without prior human review. These uses require different assessments because their consequences differ (**Figure 1**). The inventory must therefore include regulated, internally developed, purchased, embedded, general-purpose, and research systems whenever they use hospital data or

influence live work. Non-clinical uses also belong in the inventory when they affect people or access to services, even if they have no direct role in bedside care.

The hospital may already have assessed a safeguard through a medical-device, laboratory, research, privacy, cybersecurity, or enterprise-technology review. The AI Steering Committee may use that finding when it remains applicable to the proposed use, version, information flow, setting, and conditions. The registry and dossier must still identify the accountable owner, risk tier, permitted use, decision, monitoring, and next review date. They must also show how changes and incidents will be handled.

6.1 AI registry and dossier

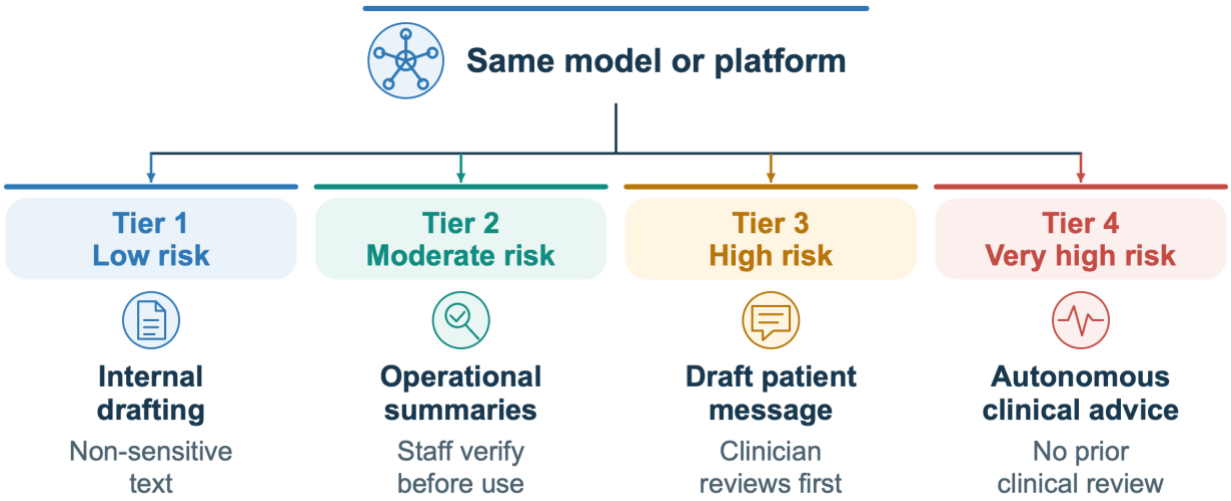

***Figure 1**. A single AI system can support uses with different credible consequences because the people, data, workflow, and downstream actions differ. The examples illustrate possible Tier 1–4 classifications after the Tier 0 screen. They do not assign tiers by task label or product category; the credible consequences of the defined use determine the tier. Each intended use requires its own scope, evidence, approval, and monitoring.*

The AI registry is the hospital's inventory of proposed, evaluated, and operational AI uses, including unapproved uses under review. Each entry identifies the service, its owners, proposed or permitted use, how it is supplied, and how it handles data. Each entry has one current status: Research only; Intake; Local assurance; Approved but not live; Early-adopter implementation; Routine operation; Restricted; Paused; or Retired. Local assurance means collecting and assessing evidence that the system and its safeguards suit the hospital's intended use. Early-adopter implementation means approved live use with an initial group, before any wider adoption. The current decision stage and requested decision are recorded separately. For example, a service can remain in Routine operation while an expansion is reviewed. [15] Actual use is compared with the registry to identify overdue reviews, unsupported versions or setups, and use outside approval.

The AI dossier is the evidence and decision file for one use. It starts with the requester's account of the problem, current process, plausible non-AI alternatives, and hospital objective. Missing essential information must be clarified before reviewers assume the proposal is useful or feasible. Approval for operational use requires a credible value case, accountable owners, and a workable operating plan. The dossier preserves the evidence and specialist findings, unresolved uncertainty, and limits on permitted and prohibited activities. It links these to operating procedures, staff preparation, the final check before live use, monitoring, changes, and incidents. This history lets later reviewers understand why decisions were made and whether their basis still holds. [15] Model cards and dataset datasheets can supply structured information about intended use, performance, and the origin and handling of data. [32][33]

The registry identifies the decision currently in force and links it to the supporting detail in the dossier. Links to approved source documents are used where possible, with the versions needed to reconstruct earlier decisions preserved. The approved baseline is the recorded version, setup, intended use, and conditions against which later changes are assessed. **Appendix I** defines the common record fields, and **Appendix L** provides the adaptable record for one use. A search or conversational interface may help staff find information, but requires its own governance review if it uses patient data or influences clinical work.

6.2 Unregistered AI and approved services

An intake route will not identify every AI use. Procurement records, software and integration inventories, research portfolios, support requests, and periodic confirmation from departments can reveal activity not submitted for review. Privacy or cybersecurity staff may investigate a specific external data flow when doing so is lawful and necessary. The extent of these checks must be justified by the need to understand hospital AI use; they must not become broad staff surveillance.

When unregistered use is disclosed, reviewers first establish what is happening and contain immediate risk in a way that encourages good-faith reporting. A familiar, low-consequence task may qualify for the approved-services list after a short review. Uses involving health records, patient communication, care, access, employment, or automated action may need temporary limits until the facts are established. Reviewers must also consider the affected team's fallback. Abruptly withdrawing a service on which staff already depend can create further risk.

The approved-services list should let staff check whether their intended task is permitted. Each entry states where and by whom the service may be used, what information and tasks are allowed, what checks users must perform, and where they can obtain support. It names the service owner, review date, and route for concerns. A material change to the provider, features, data handling, system connections, or actual pattern of use requires review of the entry so the published instructions continue to match the service available.

7 Accountability and decision authority

Figure 2 shows how hospital oversight connects to decisions about individual AI uses. The board or equivalent governing body sets the kinds of risk the hospital is willing to consider and its expectations for oversight, within the hospital's lawful authority structure. Executive leadership provides resources and assigns decision-making authority within stated limits. The AI Steering Committee and named service owners make the decisions assigned to them; unresolved matters go to the authority able to decide them. A shared network may provide expertise at any level, but the local decision-maker must remain identifiable. The governing body should connect AI investment to a small number of outcomes that guide funding and review, whether clinical, operational, educational, financial, or patient-facing. The committee reports against those outcomes and refers risks beyond its authority. A high risk tier alone does not require board approval.

Executive leadership provides the staff and resources for governance and connects the work to clinical, safety, data, technology, legal, and commercial teams. It decides matters beyond the committee's authority or permitted risk, such as policy exceptions, major contractual or liability concerns, threats to critical operations, and high risk remaining after safeguards. It also determines what resources are needed to address them. Specialist findings continue to limit what the hospital may authorize.

Within its mandate, the committee assigns tiers, approves uses and conditions, and reviews monitoring and material changes. It decides whether a service may continue, requires correction or restriction, should be paused, or must be retired. Members must have the knowledge and skills to assess the evidence; a senior signature

cannot make up for an inadequate review. The chair should have a direct route to the executive leader or leadership group responsible for resolving escalated matters.

The clinical or end-user, technical, and operational owners organize the work for each service. They obtain local evidence, fit the system into the workflow, prepare users, arrange support, and coordinate responses to changes or problems. They involve specialists whenever their expertise or mandate is relevant, while remaining responsible for coordinating the service as a whole. Purchased and partner-provided tools also need an internal technical owner and a contracting contact. Supplier support does not replace the hospital's responsibility or authority to restrict use.

Hospitals that cannot staff a full local committee may obtain support through a hospital group or a regional or professional network. Sharing reviewers, methods, training, supplier requirements, and incident lessons can reduce duplicated work and fill gaps in expertise. Each hospital remains accountable for its local decision and service. It must decide whether the shared evidence and support are sufficient for its setting.

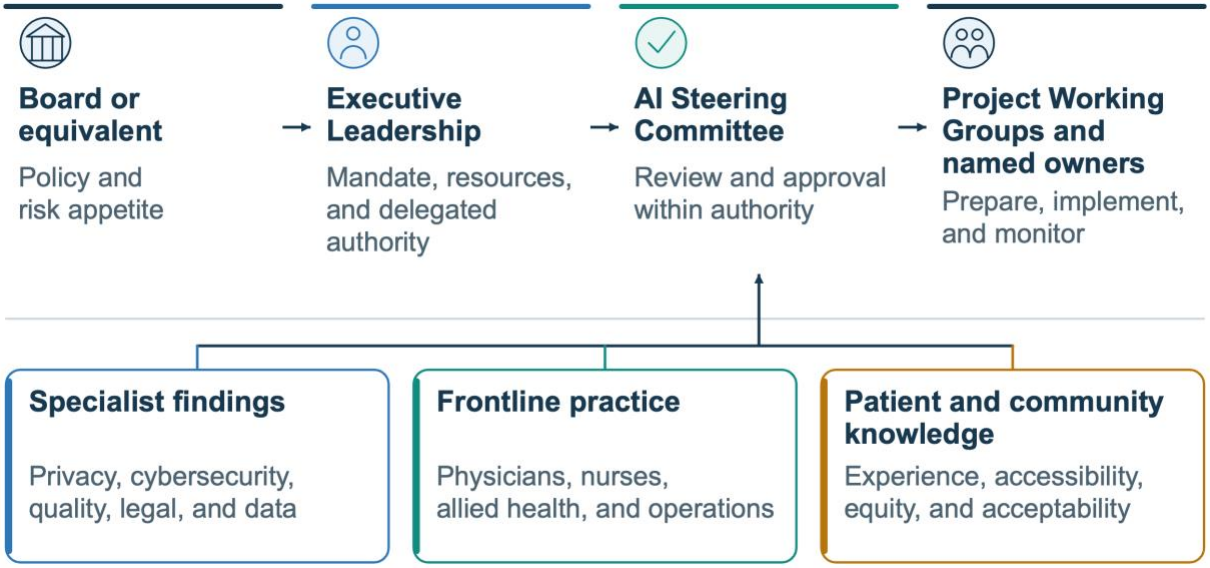

***Figure 2**. Institutional accountability connects board policy and risk appetite with executive resources and delegation, governance decisions, and implementation by Project Working Groups and the three named owners. Specialist findings and the experience of frontline staff, patients, and communities inform decisions within the authority assigned to each role.*

7.1 Responsibilities for each AI service

Every approved use needs a clinical or end-user owner, an internal technical owner, and an operational owner. Their responsibilities are recorded in the dossier. [19][14] These are responsibilities for the service and its effects, not ownership of intellectual property. The clinical or end-user owner is accountable for whether the AI is suitable for its task, how it is adopted and changed, and the resulting benefits or harms. The operational owner is responsible for making the service work in everyday practice, including workflow, training, support, fallback, and the staff and resources needed to sustain it. Because incidents and changes often require joint action, the dossier specifies how the owners pass information and responsibility between them.

The internal technical owner is accountable to the hospital for checking that the local system meets its technical requirements. This includes its approved version and setup, data quality, connections to other systems, and access to evidence and logs. During use, the owner interprets monitoring findings, assesses changes, and coordinates technical investigation. The owner also arranges restriction, escalation, or rollback when needed. Rollback means restoring an earlier approved version, setup, or service state. If rollback is unavailable, the

hospital requires an approved, tested plan to disable the affected service and continue safely. The hospital does not have to perform all engineering work itself. The hospital performs work on supplier-controlled components only when the responsibility agreement assigns that work to it.

7.2 Technical responsibility for purchased AI

Purchased and partner-provided tools still need an internal technical owner. A supplier may control design, hosting, updates, and repairs, but the hospital must judge whether the service is safe and suitable in its own setting. The owner determines whether the hospital can test and monitor the service, interprets signs of worsening performance, and coordinates investigation. This includes questioning unsupported supplier claims and referring concerns to the appropriate authority when they cannot be resolved within the approved arrangements.

A hospital that mainly buys AI can obtain technical assessment from shared reviewers, an independent assessor, or a regional service. The agreement states who will do the analysis, while the hospital remains accountable for the local decision. It must retain enough knowledge and evidence to judge whether the service remains safe and appropriate for its patients, workflow, data environment, and permitted level of risk.

The dossier and contract include a component responsibility matrix—a table of duties for each part of the service. For every material duty, including monitoring and incident response, it states who does the work, who checks the evidence for the hospital, and who may decide what happens next. It specifies the response deadline and what happens if the duty is not fulfilled. The arrangement must give the hospital enough evidence from its own systems and the supplier to investigate concerns and act. If an essential duty has no responsible party or cannot be verified, the hospital must revise the contract, narrow the use, hold the affected activity, or stop it, as the unresolved requirement demands.

Even when technical work is performed by a supplier or partner, the hospital retains authority over the limits of approved use and acceptance of the risk remaining after safeguards. It also retains the decisions governing the start of live use, expansion, restrictions, pauses, restart, and exit. For services that can materially affect care, access, rights, patient communication, or the official record, these governance arrangements must explain how patients, families, clinicians, and communities can raise concerns and request explanations. Where appropriate, they should also offer opportunities to participate in oversight. [6][21][22]

7.3 Committee expertise and participation

The committee must have access to the expertise needed for each proposal. This includes accountable executive or senior administrative leadership, clinical leadership and nursing or clinical-operations leadership, quality and patient safety, and technical expertise in data science, informatics, analytics, or machine learning. Cybersecurity, IT, privacy, legal, procurement, and equity expertise must also be available. The minimum model may use part-time members and bring in specialists when their duties apply. Membership and staffing can reflect the hospital's AI uses and existing organization, but every required finding must come from the team authorized to make it.

People whose work, rights, or decision-making authority are affected must be involved early enough to influence both design and decision. [38][39][40] Clinical, pharmacy, allied-health, and diagnostic teams can identify dependencies on other people, systems, or resources and the points where information or responsibility passes between teams. Learners, administrative teams, and support staff may reveal burdens, workarounds, or losses of capability missed by the project team. Patients, families, and communities contribute evidence about care, communication, and access. Relevant Indigenous representatives exercise their governance authority as well as contributing knowledge. Specialists participate when the use falls within their mandate, whether it

concerns human factors and education, ethics and equity, accessibility, health economics, research, professional practice, employment, or occupational health and safety. The decision record must preserve these findings and identify those that limit what the committee may authorize.

7.4 Decision authority and escalation

Every hospital must name who is accountable and who may make each AI decision, including when existing staff perform the work part-time. The hospital maintains an authority matrix stating which role may make each decision, the limits of that authority, and where matters beyond those limits must be referred. The comprehensive model adds dedicated staff and regular coordination. Under either model, the risk tier determines the evidence and oversight needed; the authority matrix identifies who may decide whether the remaining risk is acceptable.

A properly constituted AI Steering Committee acting within its formal mandate may approve a Tier 3 or Tier 4 use only when the evidence, specialist findings, and safeguard checks due for that decision are complete and the remaining risk is within its authority to accept. For each condition due before live use, the record states which activity must wait and who is authorized to verify that the condition has been met. The decision requires referral to executive leadership or the board when the committee lacks authority or the remaining risk exceeds its permitted limits, including serious unresolved safety or equity concerns and substantial risk to the hospital.

7.5 Approval and production confirmation

The AI Steering Committee coordinates review, while teams with independent authority under law, professional requirements, or hospital policy retain that authority. When applicable, privacy, cybersecurity, clinical safety, research ethics, legal, records, procurement, and workforce findings must be signed or linked in the dossier. A committee vote cannot override a mandatory finding. An unresolved requirement continues to block or limit the activity to which it applies.

The approval states the precise use and version, what is permitted and prohibited, the evidence considered, and the risk accepted. Every approval condition identifies the responsible person, due date, decision or activity that must wait, and person authorized to confirm completion. The approval has an expiry or review date and includes acknowledgement of responsibilities from the three service owners. If a condition for live use remains unresolved, the approval allows only the preparatory work explicitly identified in the decision.

Production confirmation is the final decision before live use. The designated authority checks that the actual version and setup meet the approval, including data routes, integrations, prepared users, monitoring, support, and fallback. These checks must take place in an authorized environment without enabling unapproved patient-facing activity. Every requirement due before live use must be met and verified. If the actual service differs from the approval, the hospital must correct the difference, return the service for focused review, or keep it out of live use. The affected part of an approved expansion, a material change to the live system, or a restart requires another check.

7.6 Research into routine use

Research systems may use hospital data, infrastructure, or workflows before they are ready for routine service. These systems must be registered whenever they use those resources or could plausibly become a routine service, so dependencies and responsibilities can be traced. Study permission does not authorize reliance on the system in ordinary work. The registry distinguishes research-only systems, silent evaluations using live data or workflows, and operational services that affect care or hospital work. [34][3][20][35]

A silent evaluation tests the system using live workflows or data while preventing its outputs from guiding care or operations. It requires an explicit prohibition on reliance, named owners, proportionate privacy and security review, and a route for workflow or safety concerns. These requirements apply even if end users cannot see the outputs.

Research ethics approval addresses the conduct of a study; it does not establish readiness for routine operational use. SPIRIT-AI and CONSORT-AI support transparent reporting of AI trial protocols and results, including the intervention, its setting, and human–AI interaction. [36][37] Before routine use, the hospital must assess the version and setup proposed for routine use in the intended population, with evidence that the technology, workflow, monitoring, response arrangements, and fallback are ready. Study findings may contribute to this assessment, but the hospital must still complete operational intake and the applicable reviews before production confirmation, the final authorized check before live use. [15][3][20]

The research-to-operations section of the dossier records the purpose proposed for routine use, the authority for using data for that purpose, and any patient, Indigenous, or community requirements that apply. It names the three owners accepting operational responsibility and shows how support and monitoring will be sustained under ordinary conditions. Where a supplier or partner remains involved, it assigns duties for material changes, incidents, continuity, and exit. The functions responsible for intellectual property, conflicts, regulation, credentials, and other mandatory matters retain authority over their findings. Research and operational governance may reuse relevant evidence, but approval for one does not replace approval for the other.

7.7 Starting and expanding services

The registry, dossier, and companion record use the nine statuses in **Section 6** consistently. Approved but not live means governance approval has been granted but production confirmation or another prerequisite for live use remains outstanding. Early-adopter implementation and Routine operation both mean authorized live use. Restricted means use is limited, Paused means it is suspended, and Retired means it has ended and the required closure work is complete. An expansion is a requested decision within stage 6, not a status. A vendor proof of concept (POC) is a time-limited evaluation to inform purchasing, with its own permissions and limits. Before any live use, the hospital must confirm ongoing funding, named owners, monitoring, support, and a plan for ending the service, and maintain these arrangements during operation.

When a new service needs closer initial support, early adopters should work in an area that has a measured need and sufficiently representative work to inform later decisions. That area needs capable support and practical ways to report problems and contain harm. The record identifies the approved version and setup, the maximum extent of this initial use, what users must demonstrate in practice, the date of the first review, and who can stop the service. Reviewers must examine whether safeguards work in practice; an instruction to be careful is not evidence that they do.

A proposed expansion needs a new decision if it changes the approved population, setting, volume of work, or capabilities. The existing service may continue within its current approval while that decision is reviewed. Reviewers compare evidence from current use with the needs and conditions of the receiving site or team, including staffing, monitoring, supplier support, fallback, and ongoing funding. The expansion decision must specify what is added, and production confirmation must be completed for that added part before live use begins. Review of any established service begins with its existing approval and operating history. Current evidence then informs whether it may continue, needs correction or restriction, can expand, or should be retired.

8 Sourcing, procurement, and change control

A hospital can develop an AI service itself, buy a product or managed service, or work with a partner. The choice determines who performs technical work and which duties must be secured by agreement or contract. In every case, the hospital remains responsible for approval and safe use in its own setting. [3][4][20][14] Before committing, the sourcing decision must establish whether the arrangement provides the evidence and control needed to validate the service and agree terms for its use. It must also support ongoing monitoring, incident response, continuity, and eventual retirement. [15][19] **Figure 3** distinguishes the hospital's continuing responsibilities from work assigned to suppliers or partners.

Assessment considers where and why the service is used, as well as what its AI does. Clinical care, administration, and research are areas of use; prediction, content generation, and automated actions are capabilities. Features such as recording conversations, retaining information between sessions, accessing other software, or adapting over time may require additional expertise and checks. These descriptions help assign reviewers. The risk tier depends on the consequences of the particular use, rather than its category or product name.

8.1 Procurement standards for AI tools

Procurement must secure the evidence, access, remedies, and exit rights the hospital will need during use. Reviewers must judge whether those rights are sufficient to evaluate, oversee, and control the proposed version and setup in the intended workflow. Functionality, price, reputation, and availability throughout the hospital do not answer those questions. Data governance determines who may use information, whether it is suitable, who is responsible for it, and how it is controlled. AI review then examines how the service uses that information to affect decisions or actions. Permission to use a dataset does not automatically cover a new supplier, purpose, software connection, or decision. Health-data guidance from the World Health Organization's Regional Office for Europe (WHO/Europe) and the Organisation for Economic Co-operation and Development (OECD) informs this assessment. [58][59]

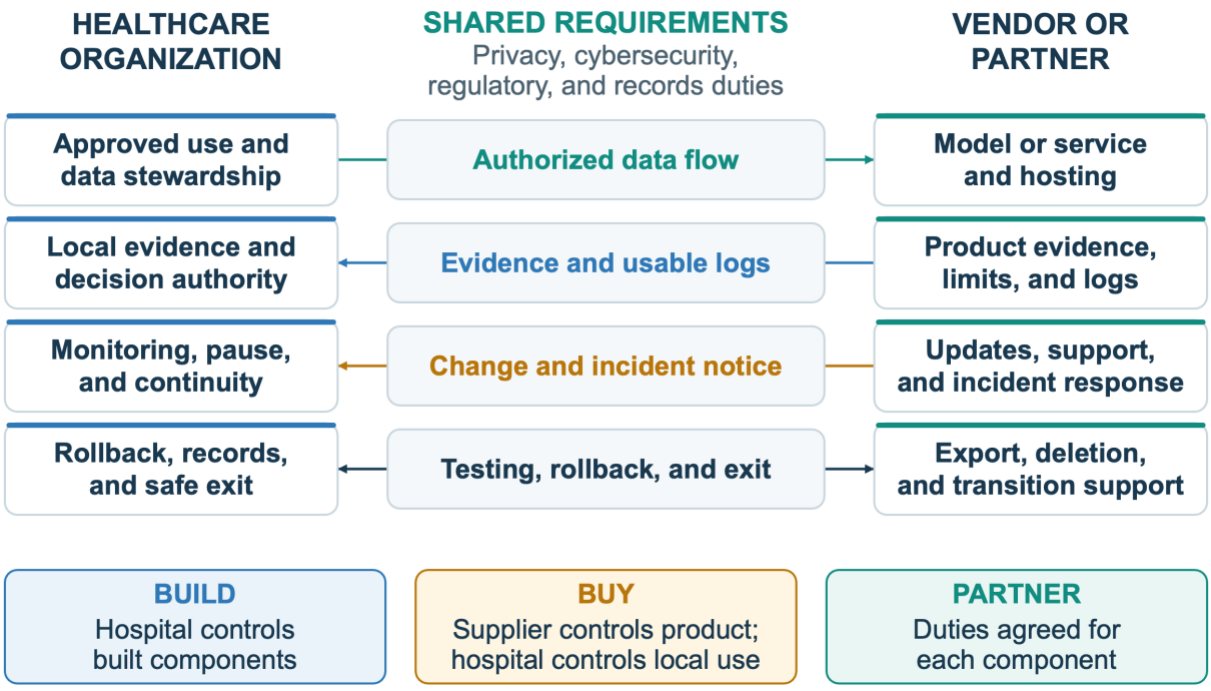

***Figure 3**. Build, buy, and partner arrangements allocate development and support work differently. Local oversight depends on access to data, evidence, information about changes, and testing. The hospital retains accountability for approval, production release, monitoring, continuity, and retirement of its use under every supplier arrangement.*

Cybersecurity assessment should begin with the hospital's established controls for identity, applications, infrastructure, suppliers, updates, and incident response. It should then examine threats introduced or changed by the AI system's design and connections. Generative and action-taking services may introduce prompt injection, unsafe tool use, or source material deliberately altered to mislead the system. They may also create logs containing sensitive information or rely on components or services absent from conventional software. The assessment must establish whether the hospital and supplier can detect and contain these threats, investigate what occurred, and recover while maintaining a safe alternative. [60][61][62]

At minimum, procurement reviewers must establish the permitted and prohibited uses, supporting evidence, actual information flow, ongoing support, and update arrangements. They must determine whether the hospital can monitor the service and investigate problems. Procurement may proceed only if the hospital can require corrections, restrict or disable use, preserve necessary data rights, maintain continuity, and leave the arrangement safely. A comparative supplier scorecard can help show trade-offs, but a high overall score cannot compensate for an unmet requirement.

Reviewers must first establish whether the evidence supports the proposed use and whether the system fits the intended workflow. They must also assess whether the hospital can sustain acceptable usability, equity, data protection, supplier support, lifecycle cost, and exit arrangements. This includes interoperability—the ability to exchange and use information with other systems. Discounts, grants, and in-kind contributions are disclosed separately from recurring costs so that temporary commercial support does not obscure the long-term value case. **Appendix G** provides detailed checks for procurement and for data leaving the service; **Appendix L** records the decision and links to the supporting evidence.

The contract must make the agreed responsibilities enforceable. It must name hospital and supplier contacts, assign monitoring and incident duties, and secure the evidence, access, cooperation, and response times needed to investigate and act. It must also provide enough advance notice and evidence to assess changes to the model

or product, hosting, data handling, interfaces, setup, or intended use. If an obligation is not met or risk becomes unacceptable, the hospital must be able to require correction, defer or disable a feature, restore an earlier version, restrict use, or exit while preserving necessary data and continuity. For a consequential service, essential evidence and remedies must be contractual rights.

8.2 Vendor proof of concept

A vendor proof of concept (POC) is a time-limited evaluation used to answer a purchase or configuration question before committing to a product or setup. Under this framework, a separate POC is required when existing evidence or an equivalent authorized evaluation cannot resolve an important question about suitability, performance, integration, or workflow fit. If it is omitted, the procurement owner and governance authority must record why and cite the evidence, following local policy. An internal build without vendor selection does not need a vendor POC, but every use still needs proportionate testing and review in the hospital. The POC rule is a proposed governance policy. Applicable legal duties must be assessed separately.

Before a POC begins, procurement must specify the purchasing question it must answer, the products or setups being compared, the comparison option, and test scenarios. The procurement owner, evaluation lead, required reviewers, evaluators, and decision-maker must be named. The authorization must set the schedule, cost, expiry, acceptance criteria, reasons to stop, incident-reporting route, and arrangements for ending access and disposing of information.

The authorization must specify the environment and limits on data and access. Artificially generated or de-identified data and connections to test systems rather than operational systems are preferred. Use of personal health information (PHI), contact with live systems, or patient-facing activity requires the same specialist review as outside a POC and separate authorization before work begins. A live pilot also requires the applicable limited operational approval and production confirmation, because procurement evaluation alone cannot authorize operational reliance.

During the POC, evaluators must follow the authorized protocol and record results as work proceeds. The record links every scenario and setup to its findings, important errors, supplier interventions, deviations, and evidence. It distinguishes results from the original product from those following a change made during testing. Vendor comparisons use the same need, scenarios, measures, and conditions, with unavoidable differences recorded. The pre-set criteria apply even if a product is free or the hospital has already invested heavily in integration. Vendor-selected evidence and evaluator conflicts must be disclosed. Vendor staff may assist but cannot select cases, alter results, or make the purchasing decision. Changes to the product, setup, data, access, or criteria need an approved amendment. The affected work must stop when a stop condition or suspected safety, privacy, or security event occurs.

After testing, reviewers must compare results with every pre-set criterion. Their assessment explains material errors, workflow burden, interoperability limitations, support needs, cost, and exit implications. It identifies differences between what was tested and what would operate in the hospital. The authorized decision-maker may stop procurement, request more evidence, revise the specification, compare another option, negotiate conditions, or continue. Every unresolved condition retains an owner and due date. The findings inform purchasing, contracting, and local assurance; they do not grant governance approval or confirm production readiness. The purchasing decision does not authorize production data access, operational use, release, or expansion.

At the end of the POC, temporary access, connections, and credentials are removed. Closure includes verification that data, copies, logs, and outputs have been retained, transferred, or deleted as authorized, with responsibility assigned for any unresolved incident, cost, contractual issue, or evidence gap. If procurement continues, contracting and local assurance require their own authorization; POC access, acceptance criteria, and

permissions do not carry forward automatically. **Appendix G** gives the detailed checks, and **Appendix L Section 5.2** records whether a separate POC is needed and its findings or the reason for omission. Calling implementation, early-adopter operation, or expansion a POC does not remove the requirements for live use.

8.3 Build, buy, and partner arrangements

Under every sourcing route, the hospital retains authority over local approval, the start of live use, expansion, continuity, restriction, restart, exit, and retirement. This remains so when another party controls development, hosting, updates, or defect correction. Responsibility must be allocated by component because a service's commercial description may conceal who controls each part. A purchased service may combine a foundation model trained for broad use across tasks, locally developed software connecting it to hospital systems, and a jointly maintained knowledge source.

For each part of the service, the component responsibility matrix assigns every material duty to one accountable party. It states the evidence and response deadlines needed for oversight, reflecting what each party actually controls. The agreement must provide enforceable rights to correction, continuity, and exit, with rollback where available or an approved, tested plan for safely disabling the service. It must be updated when responsibilities change. The responsibility matrix requires the same update. An arrangement is unacceptable if the responsible parties cannot detect, investigate, correct, or contain a plausible failure.

Internal development is appropriate when the hospital needs control over design, integration, or local adaptation that a commercial product cannot provide. [3][20] That control brings a continuing obligation to maintain the service, with the clinical, technical, operational, security, and quality capacity required for its use. The build decision must therefore account for maintenance and response duties throughout the service's life as well as the initial development work.

Hospitals that do not maintain a large internal development function may rely primarily on purchased tools. [4][14] A supplier may train the model, host the service, and issue updates, while the hospital must still determine whether the exact local use is justified and safe. This requires privacy and cybersecurity review, evidence appraisal, workflow testing, and user preparation. Independent monitoring, incident response, and change control remain part of the hospital's work. Procurement decisions must connect to the operational governance record so that these responsibilities and conditions are established before use. [19][14]

A hospital that mainly buys AI needs enough technical and clinical literacy to challenge supplier evidence and judge its local applicability. The contract must make the supplier answerable for the claims, hosted service, release information, defects, and corrective work it controls. The hospital remains answerable for the approved local configuration and workflow. A product cannot be accepted without basic information about data flows and subprocessors—other organizations the supplier uses to process data. The same restriction applies when information about updates, performance limits, monitoring, or incident responsibilities is missing. Without it, the hospital cannot oversee the use.

A partnership shares work such as design, testing, implementation, or evaluation while retaining the hospital's local responsibilities. It needs an explicit agreement about who may make each decision. [3][4][14] Partners may be academic groups, startups, vendors, regional networks, or multicentre health-system initiatives. The agreement must identify who approves validation and updates, how evidence and data are obtained, and who responds to incidents. It must also assign continuity and retirement duties if the partnership ends. Intellectual-property and publication terms must preserve the hospital's ability to investigate, continue safely, and exit. Shared expertise may reduce duplication when these responsibilities are clear in routine work and during problems.

The sourcing route should reflect what the hospital can sustain. A smaller hospital may rely mainly on procurement, supplier management, clinical and technical literacy, and operational ownership. A larger hospital may also maintain staff and infrastructure to develop models and support them throughout use. A managed service, in which a supplier operates the service for the hospital, is treated as buy unless the parties share authority over design or decisions. For mixed services, responsibility is allocated by component, because the label 'hybrid' does not specify who does the work.

In every case, the record must identify who performs and who reviews each material duty. When a hospital or collaborator modifies a component that shapes behaviour, it takes on the corresponding Build or Partner responsibilities for the work it controls. Changes to decision logic, knowledge sources, safeguards, or integrations require updates to the component matrix and affected assurance and lifecycle duties. The legal and regulatory role review in **Appendix J6** must be reopened where necessary. Purchasing the original product does not remove responsibility for a local modification or its effects.

8.4 Responsibility for changes and problems

Changes in data, software, patients, or working practices can make a service depart from the conditions on which it was approved. This is referred to here as drift, and it may occur without an apparent decline in model performance. An undisclosed supplier update may change behaviour while the hospital's workflow stays the same; an unchanged product may become unsuitable after clinical practice or staffing changes. The affected part of the service must be identified so the responsible party can investigate and determine whether correction, restriction, rollback, new validation, or exit from the service is needed.

The monitoring plan must identify the approved baseline and local signals of meaningful departure, with named responsibility for investigation and response at warning, action, and stop levels. **Appendix I** provides the recording fields. Supplier reports can contribute evidence, but the hospital needs its own view of local performance, data conditions, workflow, and user practice. The supplier's account alone cannot establish whether continued use is justified.

Responsibility for investigation and correction follows who controls the affected part of the service. For hospital-controlled components, the internal team must preserve the development and change history, data records, and evidence for each release, and retain the ability to investigate, correct, and maintain continuity through rollback or an approved, tested plan for safely disabling the service. Suppliers remain responsible for the components and changes they control. The hospital remains responsible for its setup, workflow, monitoring judgments, and decisions to expand, restrict, pause, or stop. The hospital retains authority over local release. Contracts must secure the information, cooperation, notice, remedies, and exit rights needed to act. Partnerships must assign one accountable party to each component and response duty, even when work is shared. They also need a dispute process and a plan for ending the partnership. Disagreement must not leave an essential duty unassigned, especially during an incident.

Every sourcing route requires a record of the approved version, setup, and use, with its change history and controlled access, together with records that allow data flows and user practices to be checked. Material changes must be introduced in stages. Before a changed service is released, representative regression testing—retesting functions after a change—must establish that previously accepted functions and safeguards still work; incident cooperation and continuity must also be tested, including rollback where available or the approved plan for safely disabling the service. Loss of essential monitoring must trigger the approved response, including immediate suspension when a stop threshold is met. Any temporary continuity exception must follow the separately authorized emergency decision and equivalent-protection requirements in **Section 10**.

8.5 Limits on supported configurations

Different setups of the same AI use may sometimes be evaluated, monitored, and supported together within documented approved operating limits. This framework calls them a configuration group. Evidence must show that the differences do not invalidate joint evaluation or oversight; using the same product, site, or version is not enough. Each additional group creates separate work for testing, training, monitoring, support, incident response, rollback or safe disabling, and retirement. To keep that work manageable, the framework proposes a default of one configuration group per service. This is a planning recommendation, not a threshold established by research. A hospital may approve more groups if it records the justification and funds the additional work.

The decision to group setups requires comparison of the model, prompts, templates, thresholds, populations, languages, and settings. The comparison also covers data connections, possible outputs and actions, autonomy, and the workflow for checking outputs and actions. The record identifies which differences are allowed, why the evidence supports assessing the setups together, and which changes or findings require a new assessment. A material difference outside that evidence requires a separate group or revised approved limits before use. Every authorized version within each group remains identifiable.

Approved operating limits state the permitted population and setting, data and interfaces, permissions and versions, and workload. A proposed difference outside those limits needs focused change review or a new intake before use. If the hospital has reached its approved number of configuration groups, it must consolidate or retire a group, decline the variation, or fund and approve a higher limit. A temporary exception needs the same justification for adequate resources, a named owner, and an expiry date. It cannot waive a core safeguard or bypass a mandatory authority. **Appendix I** defines the record, **Appendix E** gives the checklist, and **Appendix L Sections 5, 8, and 10** record the decisions.

8.6 External information sharing

Data sharing outside the hospital requires review of who can access information, where it is processed, and what the hospital can investigate or stop. A current map must show information entering and leaving the AI service, including patient information, prompts, outputs, logs, and derived data handled by suppliers or their subprocessors. This is the ingress and egress map. It specifies the jurisdictions involved and responsibilities for access, retention, and deletion so privacy, cybersecurity, procurement, and incident reviewers can assess the arrangement.

Every use receives privacy and security triage to determine whether a privacy impact assessment, security assessment, or another locally required data or AI impact review is needed. Where a review is required, the dossier must record its scope, findings, conditions, approving authority, decision date, and next review or expiry. Those findings apply to the assessed version, data path, and supplier arrangement. A material change requires review before they can be assumed to remain applicable.

The approved system design must minimize exposure and restrict access while preserving the evidence needed for investigation. Enforceable terms must govern notification, retention, deletion, and any use of information beyond the original purpose. Whenever the risks of access or processing warrant it, hospitals using either governance model must review subprocessors and jurisdictions, verify actual data flows, and check supplier compliance. Dedicated governance teams may conduct these audits routinely; **Appendix G** records the detailed evidence for each use.

Elevated permissions, remote access, cloud access, and access to sensitive data require strong authentication. Each service identity—the account or credential software uses to identify itself and gain access—must have a named owner and only the permissions it needs. Secret credentials must be stored securely, with defined expiry and replacement arrangements for API keys and other service credentials under hospital security policy.

Compromised or unnecessary access must be revoked promptly. Password policy should follow current authentication guidance, which does not support arbitrary periodic changes in the absence of compromise. [135] Any exception to required protection needs approval, a time limit, and a tested alternative that provides equivalent protection.

8.7 Updates and material changes

Under every sourcing route, a change requires review if it could alter the approved purpose, data path, performance, integration, human oversight, patient-facing use, supplier dependence, or credible risk. Retraining, replacement of a model on which the service depends, changes to prompts or guardrails, and movement to a new setting may meet this test. Guardrails are rules or mechanisms that limit what the system may produce or do. As with other changes, reassessment depends on the effect on approval rather than the technical label.

Both governance models require a documented change assessment, testing proportionate to the change, a specified approval authority, arrangements for rollback or safe disabling, and review after release. A hospital using the minimum model may use a concise change request with version and approval records; a hospital using the comprehensive model may automate routing and monitoring. Both must preserve the evidence and decisions. [16][17][18] The internal technical owner investigates local signals and involves suppliers or specialists when their evidence or action is needed.

Unresolved or repeated problems require referral to the AI Steering Committee, which decides whether the service can continue, must be restricted or rolled back, needs renewed validation, or should be retired. Executive or board review is required when risk exceeds delegated tolerance, a mandatory concern remains unresolved, or the contract is unenforceable. The same requirement applies when loss of a service essential to hospital operations creates substantial risk to the hospital. A referral must identify the authority needed to resolve the issue; escalation does not permit a mandatory concern to be set aside.

9 Risk tiering and evidence for approval

The framework draws on risk-based approaches in the Algorithmic Impact Assessment, NIST AI Risk Management Framework, and EU AI Act. [41][42][43] Its tiers describe the most serious harm that could plausibly follow from a particular AI use. Reviewers consider what the system can do within safeguards that are already implemented and verified, including any required human approval before an action. Planned or untested safeguards cannot justify a lower tier. Reviewers then make a separate judgment about whether the evidence and functioning safeguards make the remaining uncertainty and risk acceptable within their authority. A high-tier use may be approved when these requirements are met; effective safeguards do not necessarily change its tier. **Figure 4** shows this sequence.

Tier 0 is the first screen of the next proposed activity. Reject applies to a prohibited use or an unacceptable feature that cannot be corrected. Hold applies when essential information or a prerequisite is missing. Proceed states the limited review or evaluation that may start; it does not authorize live use. The required specialist reviews are identified at intake, even if their final findings will come later. Access to data, connections to live systems, and patient-facing work must each wait for the permissions required for that activity.

Figure 4. *Tier 0 screens whether the next proposed activity must be rejected, held, or allowed to proceed. Reviewers then assign the highest applicable consequence tier for the defined use within its verified limits. Evidence and tested safeguards inform the separate decision on whether the remaining risk is acceptable. Governance approval permits only the specified activity; live use also requires production confirmation.*

9.1 Risk appetite and risk tolerance

Risk appetite describes the kinds of risk a hospital is willing to consider in pursuing benefits for care, experience, equity, or service performance. Risk tolerance sets the limits on what may be accepted in a particular decision. For example, a committee may approve a Tier 3 use if the remaining risks are within the limits assigned to it. Those risks may concern safety, privacy, equity, performance, workload, or actions that are difficult to reverse. Risks or exceptions beyond those limits must go to the person or body authorized to decide them. Assigning a tier does not itself give permission to proceed.

9.2 Tier 0 screening and consequence tiers

Tier 0 records Reject, Hold, or Proceed for a specified activity without assigning a consequence tier. A hold may be needed if ownership, essential information, access to evidence, or lawful data authority is missing. It may also be needed if the proposal lacks required human control, records sufficient to reconstruct system activity, or a workable fallback. The proposal remains on hold while it is clarified or redesigned. A prohibition or an unacceptable feature that cannot be remedied requires rejection. For Proceed, the record states which activity may begin and under what conditions. Every Tier 0 decision must record the reason, any corrective action, and who may reconsider it. **Appendix D1** provides the Tier 0 checklist.

Tier 1 (low risk) applies only when failure could cause minor, readily corrected inconvenience and a user checks the output before relying on it. The use cannot affect care, clinically significant access, patient communication, rights, employment, resource allocation, or the official clinical record. Drafting non-sensitive internal material or summarizing published literature without patient-specific information may qualify within these limits. The service still requires a named owner, a concise factsheet stating its purpose and limits, and a basic workflow

check. It also needs an initial privacy and security assessment to identify any further reviews, and a route for reporting and responding to incidents.

Tier 2 (moderate risk) applies when failure could cause meaningful but limited disruption, workload, delay, or a correctable privacy, fairness, or service concern. Correction must be timely, with no credible pathway to serious harm within the verified use limits. Scheduling or capacity support with defined limits may therefore qualify when it does not determine care or clinically significant access. Review must establish the documentation, ownership, workflow, and value case for the use, together with the required privacy and security review, monitoring, and a safe alternative.

Tier 3 (high risk) applies when an output can materially influence care, access, employment, rights, resources, prioritization, patient communication, or an official record, with meaningful human review before consequential action. Examples include a clinician-reviewed deterioration alert, a signed note drafted by an AI scribe from a captured clinical conversation, or a chart summary for clinician review. A coordination action within specified limits, approved by an authorized person, may also qualify. [15][34][18] Evidence must address local technical performance and the workflow in which people interpret and act on the output. It must include human factors, equity, monitoring, and change control. Reviewers must also assess user readiness and a workable fallback. Approval depends on the human decision point functioning in practice.

Tier 4 (very high risk) applies when a system can drive or carry out a time-critical or similarly consequential decision or action without meaningful case-specific approval, or when serious harm can occur before meaningful review. Autonomous treatment, dosing, diagnostic advice, clinically significant access decisions, or final interpretation may qualify according to their actual consequences. These uses require all Tier 3 controls, supplemented by stronger and more independent evidence, tighter technical limits, and closer monitoring. Any remaining risk beyond delegated tolerance must be referred to the authority able to decide it.

A tier applies to a defined use because the surrounding service determines what harm a product can cause. Patient-facing output, sensitive information, deeper integration, external data transfer, or less opportunity to intervene can create new routes to harm. Reviewers must reassess the tier when that context changes. A model or vendor therefore does not have a permanent classification independent of its use.

9.3 Tier selection and acceptance criteria

After Tier 0, consequence tiering starts at Tier 4 and proceeds downward until the first applicable tier is reached. The decision reflects what the system can do within limits that are implemented, enforced, and verified. A credible harm pathway is a plausible sequence in which an error, failure, or foreseeable misuse leads to harm. The record includes the evidence and uncertainty supporting that judgment. Likelihood informs the assessment of plausibility, but an average result cannot rule out a credible severe outcome. Human review can limit the consequence only if it is enforced in practice and the reviewer has enough information, time, skill, authority, and a workable alternative to reject the output before harm occurs. Later assessment and monitoring must check whether that review continues to work.

1. **Use definition.** The assessment states who receives what output, in which setting and population, and from which data. It explains what decision or action the output can change, where the result is recorded or sent, and whether the system can act. Each materially different use receives a separate tier.

2. **Tier 0 screening.** The screen identifies prohibited features, missing essential information, and barriers to the next activity. It covers ownership, purpose, data authority, records sufficient to reconstruct system activity, required human control, stop and fallback capability, and supplier evidence or cooperation. The decision is recorded as Reject, Hold, or Proceed, with its specific scope and conditions. Specialist questions go to the appropriate teams, whose required findings must be available before the activity to which they apply.

3. **Harm pathway.** The assessment follows a wrong, missing, delayed, unavailable, misused, or materially changed output through the actual workflow. It records the most serious credible consequence and time to harm, the people and groups exposed, their vulnerability, and the scale of exposure. It considers how readily failure can be detected and its effects reversed. The record also covers effects on official records or patient-facing services, the service's dependence on AI, and the actions the system and users are authorized to take.

4. **Tier 4 assessment.** Can the system drive or carry out a time-critical or similarly consequential decision or action without meaningful case-specific approval by an authorized person, or can serious harm occur before meaningful review? An affirmative answer indicates Tier 4.

5. **Tier 3 assessment.** Can the output materially influence care, access, rights, employment, resources, prioritization, patient communication, an official record, or another consequential action, with meaningful human review before action? An affirmative answer indicates Tier 3.

6. **Tier 2 assessment.** Could failure cause meaningful but limited disruption, delay, workload, or a correctable privacy, fairness, or service concern, with timely correction and no credible pathway to serious harm within the verified use limits? An affirmative answer indicates Tier 2.

7. **Tier 1 assessment.** Tier 1 applies only when no higher-tier pathway is credible, consequences are minor and readily corrected, and a person checks the output before relying on it. The output must not affect care, clinically significant access, rights, employment, resource allocation, patient communication, or an official clinical record.

8. **Boundaries and uncertainty.** The highest indicated tier remains provisional while a material fact is unknown. A lower classification requires evidence that an enforced technical or workflow limit prevents the higher-consequence pathway in practice. Planned safeguards and expectations of careful use are insufficient. The record explains why the selected tier applies rather than the tier immediately above or below it. It includes any disagreement and the changes that would require reassessment.

After assigning a provisional tier from the plausible consequences within the verified limits of use, reviewers assess residual risk using the applicable safeguards and evidence. A high-consequence use may proceed only when its evidence, controls, monitoring, and decision authority meet the requirements for that consequence. [44][10] **Appendix D** specifies the evidence by tier. The final decision records whether remaining uncertainty and risk are acceptable within delegated authority.

The examples show how the decision affected can change a tier even when products appear similar. Drafting non-sensitive internal material may be Tier 1 when errors are obvious and readily reversed. Capacity forecasting or scheduling support may be Tier 2 if it informs operations without determining care or clinically significant access. A clinician-reviewed deterioration alert, ambient note, or patient message will often be Tier 3 because it can affect clinical communication or action, with an enforced human decision point. Autonomous, time-critical triage or treatment action may be Tier 4.

Each example requires assessment against the actual limits of use before a tier is assigned. An enforced restriction may remove the route to more serious harm—for example, excluding clinical information, disabling a connection to another system, or requiring effective approval for each external action. A lower tier requires verification that the restriction works and reassessment of the consequences still possible. A narrower population or keeping data inside the hospital does not, by itself, establish a lower tier.

Acceptance criteria are the results a service must achieve to be approved. They are set before evaluation begins, using an appropriate comparison such as current practice, a non-AI workflow, an existing rule, or another service. Model performance alone may not show whether the service improves the intended task. The criteria should reflect the consequences of errors, how readily they can be detected, uncertainty, effects on relevant groups, operational capacity, and the work needed for safe use. Where feasible, evaluation compares the person alone, AI alone, and the person using AI under similar information and time constraints.

Evaluation must reflect differences in staff experience and the pressures of ordinary work. Human–AI performance depends on the fit between the system, task, and workflow, and on whether users are prepared. It also depends on whether users' reliance matches the system's demonstrated capabilities and limits. Even a strong model can worsen a service by redirecting attention or review in ways that undermine the task. [16] Tests must cover plausible local changes, including a new site or patient group, changes in the prevalence of the relevant condition or in data collection and processing, and service outages. The team specifies a safe response and observes whether users can carry it out. Written instructions alone do not establish readiness.

A healthcare scoping review found reported benefits from human–AI collaboration, but most studies assessed short-term task performance rather than patient or system outcomes. [45] A broader meta-analysis found that combined teams did not, on average, outperform the better of humans or AI alone. [46] Hospitals therefore need evaluation in the intended setting, including important failures and effects on clinical judgment that an average result may conceal. DECIDE-AI supports transparent reporting of early live clinical evaluation, including performance, safety, and human factors. [34] In addition to that evaluation, the receiving service must demonstrate that it has the time and capacity to respond to outputs as required.

9.4 Streamlined review for lower-risk uses

Improved staff competence may allow faster review of lower-risk tools, although further evaluation is needed to confirm this. A streamlined route in this framework requires staff with demonstrated competence, reliable records, satisfactory audits, and previously approved safeguards that still apply. Before introducing it, the hospital must set eligibility rules and establish an ongoing audit process. Reviewers compare each later proposal with the documented, approved baseline of conditions to confirm that the existing evidence applies.

Only clearly Tier 1 or Tier 2 uses may enter a locally approved streamlined route. The proposed version, setup, information flow, workflow, and safeguards must match an approved baseline. Tier 3 and Tier 4 uses are excluded. Uses that materially affect care, clinically significant access, rights, employment, an official record, or a safety-critical workflow are also excluded, as is any use with unresolved specialist concerns. Acceptance requires verification of the required registry and dossier information and confirmation that the relevant staff have current evidence of competence.

Before streamlined review begins, the AI Steering Committee approves a sampling plan and audits baseline cases that completed full review. The plan identifies the eligible Tier 1 and Tier 2 decisions, any use groupings, sampling method, reviewer, schedule, definition of a material error, and acceptance criterion. The baseline audit must meet that criterion before the pathway opens. Audits assess the pathway's reliability; they cannot extend it to higher tiers. **Appendix B** defines the recording fields.

A full review assesses all evidence relevant to the decision. A focused review examines a particular change or issue and checks that the rest of the approval still applies. Evidence may be reused only when a documented comparison shows it is current and applies to the proposed setup and circumstances for the claim being assessed. These review methods may be used within an approved streamlined route for lower-risk uses. They must preserve specialist authority and the required records for ownership, tier, monitoring, incidents, changes, and retirement.

Ongoing audits of the streamlined pathway follow the approved schedule and report findings by tier and use grouping. The committee uses these findings to judge whether the pathway still meets its acceptance criterion. The plan identifies which findings require suspension and who may order it. Before reopening, evidence must show that corrective actions have been completed and verified, staff still meet the relevant competency requirements, and a follow-up audit meets the pre-specified criterion. **Appendix D** states the evidence rule, and **Appendix L Sections 4 and 7.3** record eligibility and audit decisions.

9.5 Required specialist reviews

Every use receives initial privacy and security triage, and patient-safety review is required whenever the use can affect care. A fuller specialist assessment is required for substantial risks, including patient information leaving an approved internal path, patient communication, effects on access or treatment, or reliance on an external supplier. Review depth reflects the risk and its consequences. Duties that apply regardless of tier remain binding for the activity concerned.

Features that change how harm could occur also require focused examination in the actual workflow. Reviewers must examine ambient capture, including recording clinical conversations, content generation, and retrieval of information from the clinical record. They must also examine system adaptation, hosting outside the hospital, employment effects, and permissions to write, send, schedule, order, or call another tool. Higher-consequence uses require realistic scenarios, tested human authority, a safe alternative, and evidence that the production service meets its approved conditions.

10 Monitoring and lifecycle decisions

Approval applies to a defined use at a particular time. The data, staff, workflow, supplier, or local need may later change. **Table 1** sets out seven decision stages, covering the initial proposal, the start of live use, continued operation, changes, and retirement. It identifies the responsible authority and the record needed at each stage. [15][19][20][14] **Figure 5** groups the work supporting those decisions into six recurring tasks. After an incident or material change, the hospital may need to repeat the affected work to determine whether the approval still holds. The record identifies which decision is being reconsidered, what evidence or conditions need review, and which activities may continue in the meantime.

Governance approval answers whether the specified use is acceptable within the hospital's risk limits and under stated conditions. Production confirmation checks that the service about to go live actually meets those conditions. Confirmation requires verification that the approved model and setup are present, integrations and data routes behave as tested, users and support teams are prepared, monitoring is active, and a workable fallback is available. Every condition due before live use must be met and verified. Approval of a proposal or supplier does not authorize an unchecked later version or setup. Once the service is operating, the named owners remain responsible for referring incidents and material changes to the committee.

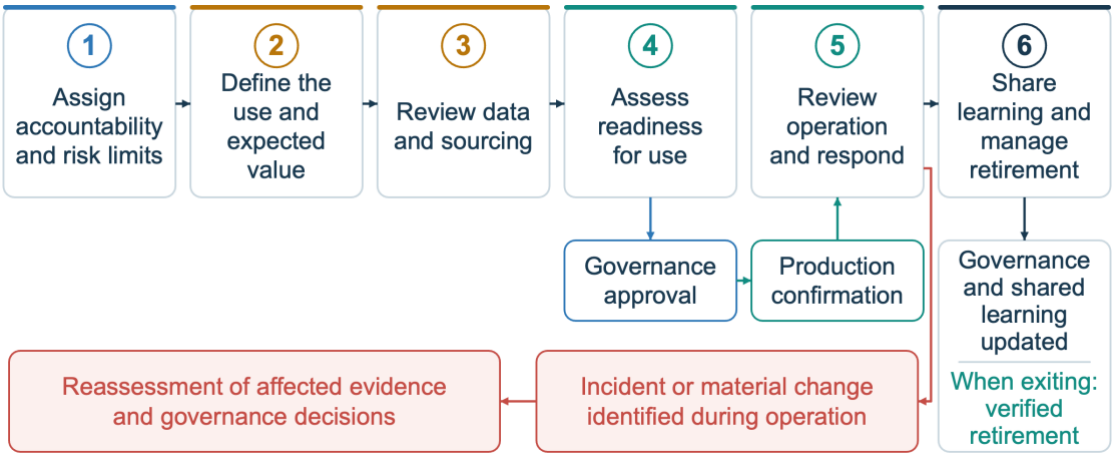

***Figure 5**. Six recurring governance tasks support the seven decision stages in **Table 1**. Governance approval and production confirmation are separate decisions completed before live operation. Incidents and material changes prompt reassessment of the affected evidence and decisions; lessons from operation and evidence of completed retirement remain in the service dossier.*

A service may need closer support when first introduced because of its possible consequences, scale, how much hospital operations depend on it, or uncertain fit with everyday work. In that case, authorized live use begins with a defined early-adopter group and limits on use. The area needs a documented need, sufficiently representative work, capable leadership, and ways to identify problems, report them promptly, and pause use safely. Initial operation tests the service and the hospital's ability to support it. Expansion requires the agreed results and ongoing clinical, technical, educational, and financial resources.

Monitoring must establish whether the operating service still fulfils its approved purpose and provides the expected benefit. It must also check that safeguards remain in place and use stays within approved conditions. When performance falls short, owners must investigate the whole service. The cause may involve the model, data quality or availability, workflow, user preparation, adoption, equity, supplier performance, or a change in clinical context. Measures should help distinguish these explanations. Each material finding should be linked to the response needed to correct, restrict, or otherwise reconsider use.

The minimum and comprehensive models require the same decisions throughout a service's life. Hospitals overseeing fewer uses may rely on controlled records, version histories, and scheduled reviews. Those overseeing more uses may need automated information feeds, software to assign reviews, and a combined dashboard. Both need timely evidence, named reviewers, and specified responses when a threshold is reached. The Project Working Group prepares the evidence with the relevant owners and specialists. The AI Steering Committee retains authority for tiering and approval.

Table 1. *Lifecycle governance stages, checkpoints, minimum outputs, and responsible oversight.*

Stage	Governance review	Required record	Who is responsible
1. Intake and triage	What exact use is being considered, and can it proceed to review?	The intake records the Tier 0 decision, provisional tier, owners, and initial privacy and security assessment. It also records how any current unapproved use has been contained.	AI Steering Committee or designated intake authority
2. Problem and value	Is the service need clear, and is AI a reasonable option?	The record defines the proposed use and exclusions, current practice for comparison, non-AI alternatives, expected benefits and burdens, and plausible ways failure could cause harm.	Clinical/end-user and operational owners with governance review
3. How the service will be provided and prepared	Can the hospital implement and oversee the proposed arrangement, manage changes, and leave it safely?	The sourcing plan explains the choice to build, buy, or partner, maps the data flow, assigns duties for each part of the service, and identifies the evidence, hospital testing, contracts, and change review needed.	AI Steering Committee with required privacy, security, safety, legal, procurement, and technical reviewers
4. Evidence and governance decision	Is there sufficient evidence for the proposed live use, and is the risk remaining after safeguards acceptable?	The decision records the approved version, setup, use, and conditions. It links to external and local evidence, findings for relevant groups and workflows, and remaining hazards.	Named owners and specialist reviewers; AI Steering Committee decides

Stage	Governance review	Required record	Who is responsible
5. Production confirmation before live use	Does the actual service match the approval and have the people and controls needed to operate?	Verification establishes whether the actual version and setup, procedures, staff competence, monitoring, support, fallback, and arrangements for rollback or safe disabling meet the approval.	Named production confirmation authority with three-owner sign-off
6. Continued use and changes to approval	Does current evidence support continued operation? Can a proposed change, expansion, or restart meet its defined conditions?	Evidence from monitoring, value and capacity reviews, changes, and incidents supports the decision to continue, correct, restrict, pause, expand, restart, or retire. Production confirmation follows when required.	Three owners with AI Steering Committee oversight and supplier escalation where relevant
7. Retirement or replacement	Can the service be retired or replaced without disrupting work that depends on it or leaving other unsafe dependencies?	The retirement record explains the reason, transition, and fallback. It records removal of access and integrations; required retention, transfer, or deletion of data and logs; contract exit; correction of affected records; and lessons learned.	AI Steering Committee with operational leadership

Appendix A illustrates these decisions through a deterioration alert, an ambient scribe, and an AI agent permitted to perform specified coordination tasks. **Appendix E** provides the lifecycle checks. The record in **Appendix L** follows the use through monitoring, change, and retirement, using the common fields in **Appendix I** and links to dossier evidence. It distinguishes the service's current status, such as Routine operation or Paused, from the stage under review. An established service can then be reassessed without losing its operating history.

10.1 Work supporting each decision

The six recurring tasks adapt NIST's Govern–Map–Measure–Manage structure to hospital work. [42] Each task may be needed at several decision stages. The intake record identifies the three service owners, the governance lead, and who may make the requested decision. It states the limits on the risk they may accept and who handles matters beyond those limits. Purchased or partnered services also need supplier and procurement contacts. The record defines the problem and intended use, including users, affected population, workflow, and activities that are permitted or excluded. It describes the expected benefit, considers non-AI alternatives, and sets the measures and baseline for assessing benefit.

The decision about how to provide the service requires review of data use and supplier arrangements, including the privacy and security assessments required. The arrangement, including any applicable contract, must provide the evidence, safeguards, remedies, and safe exit rights needed for the next authorized activity and, before live use, for ongoing operation. Evaluation then tests whether the proposed service meets its pre-set criteria in the hospital's setting. It covers local performance, effects on relevant groups, workflow, usability, human factors, and important failures, using silent evaluation where appropriate. Before live use, the actual version and setup and the preparedness of users, monitoring, support, and fallback must be verified.

During operation, the hospital must monitor the measures established for expected value and hazards and respond to the findings. A warning requires investigation, while an action threshold requires the defined action

to reduce risk or restrict use. Reaching a stop threshold requires immediate suspension of the affected function, activation of fallback, and incident escalation. These findings inform decisions to continue, correct, restrict, revalidate, replace, or retire the service. When new evidence or changed conditions may affect an approval, the relevant work must be repeated to determine whether its grounds still hold.

Findings from implementation, incidents, near misses, and retirement inform governance, procurement, and education. De-identified lessons may be shared through an authorized route. Ending a service requires verification that access, dependencies, data, and contractual duties have been addressed. Specialist teams retain authority over their findings. Each finding or obligation is linked to the use it concerns and the activity it permits or prevents, so later reviewers can identify the dependent decisions. **Appendix J** explains how these tasks relate to NIST and other healthcare guidance. [38][41][42][47][9][16][17][10][8]

10.2 Hospital-wide AI monitoring

A hospital-wide dashboard should show whether each live AI service remains within its approval, works safely and reliably, and provides enough benefit to justify its burden, cost, and risk. The registry identifies its current status, owners, and approved baseline. The dossier explains how monitoring measures benefit and plausible harms. The fields in **Appendix I** record each measure's population, data source, comparison baseline, and denominator—the total number of relevant people, cases, or events used to calculate a rate. They specify how missing or delayed data and uncertainty will be handled, how often results are reviewed, and by whom. The record must link each result to the required response and escalation authority, available fallback, and evidence needed before any restart.

Monitoring covers results and the working conditions that help explain them. The six domains cover safety and quality; technical and data performance; equity, access, and experience; user practice and workload; reliability and cybersecurity; and adoption, value, and supplier performance. Together, they show whether the service works, who benefits or is harmed, and whether staff can support it. An N/A designation for a domain or measure requires a reason accepted by the reviewer. Every applicable measure remains required.

Warning, action, and stop thresholds must be justified for the local use, taking account of how severe harm could be and how quickly it could occur. A warning starts an investigation. An action threshold requires the specified response, such as restricting use or referring a decision to an authorized senior person. Reaching a stop threshold requires immediate suspension of the affected function, activation of fallback, and incident escalation, even if the cause is unknown. Temporary continued operation requires a separate recorded emergency decision by an authorized person and safeguards providing equivalent protection. Mandatory duties still apply. Restart requires accepted evidence of correction and completed production confirmation.

Loss of monitoring can remove the evidence on which continued operation depends. The monitoring plan must specify what happens if a critical feed, audit log, incident route, subgroup measure, or supplier report becomes unavailable. A loss that meets a stop condition requires suspension of the affected function, even when the cause is uncertain. If the plan permits limited continued operation below the stop threshold, the record must identify the temporary controls, accountable decision-maker, expiry date, and evidence needed to confirm that monitoring has been restored. Each periodic review must explain why the evidence available at that time still supports operation.

10.3 AI failure patterns and monitoring

AI failures can arise from the model, its input data, the task, or the way its output affects care. **Table 2** describes common, overlapping failure patterns and checks that can reveal them. It is not exhaustive. This framework uses drift broadly for changes from the conditions supporting approval. Input data drift concerns changed inputs;

concept drift concerns a changed relationship between inputs and outcomes. Workflow or use drift concerns how people use the system. The monitoring requirements in the table draw on the cited research and guidance.

Shortcut learning and reward hacking show why a high score may not reflect useful performance. Shortcut learning occurs when a model relies on incidental patterns that predict labels in development data but may fail when conditions change—for example, features associated with an imaging device rather than the disease. [132][136] Reward hacking involves exploiting how success is measured instead of completing the intended task. If a hypothetical referral agent were rewarded for closing referrals, it could raise that score by marking them complete without delivering them. This would exploit the completion metric; an accidental delivery failure would be a different problem. [137][138]

Before live use, the three owners must review every category in **Table 2** with the relevant specialists and link applicable hazards to the risk assessment and monitoring record in **Appendix I**. For each hazard, the record identifies the safeguard, sign of a problem, reference or comparison, sampling plan, known gaps in detection, responsible reviewer, and required response. The reviewer must accept the reason for marking a category N/A. Responses to findings follow the established warning, action, stop, and approval procedures.

***Table 2**. AI failure patterns and the checks needed to investigate them.*

Failure pattern	What can go wrong	Monitoring and investigation
Target and representation bias	The prediction target is the outcome or quantity the model is trained to estimate. Target bias can arise when this is a poor proxy for the intended benefit. Representation bias arises when the data inadequately reflect relevant people or conditions. Predicting healthcare cost accurately, for example, need not identify clinical need fairly. [84]	Review covers errors, calibration (whether predicted risks match observed outcome rates), access, and benefit by relevant group, including people excluded from scoring. It examines whether the target still represents the intended benefit and reports group sizes, missing demographic data, and uncertainty. A small sample or nonsignificant difference cannot establish equity.
Data noise and label noise	Measurement error, corrupted signals, transcription errors, and missing values can affect development, live use, and evaluation. Label noise means errors or inconsistencies in the recorded answer or outcome used as a reference. Disputed labels also require review; disagreement does not by itself show which label is wrong. [133]	Checks cover completeness, units, timing, and signal quality. Sampled records are compared with their sources, and disputed labels receive qualified independent review. Genuine disagreement is recorded. Disagreement with the model's answer alone is not a reason to replace a label.

Failure pattern	What can go wrong	Monitoring and investigation
Covariate shift and out-of-distribution inputs	Input data drift is a change in the distribution of input data, for example as patient mix, equipment, or data-collection conditions change. Covariate shift assumes that the relationship between inputs and the outcome being predicted stays the same; reviewers must justify that assumption. Out-of-distribution inputs come from a different data distribution than the one used to develop or evaluate the system. [130]	Inputs, missing data, devices, sites, and populations are compared with the approved baseline. Checks include outcomes and errors when available and local evaluation of unfamiliar-input detectors. A drift alarm warrants investigation; no alarm does not establish safety. [131]
Prevalence changes and label shift	The outcome may become more or less common. The term label shift assumes that input patterns within each outcome category stay the same. Changes in coding or outcome recording can look like changes in prevalence, even when the underlying outcome rate has not changed. [130]	Monitoring tracks verified outcome rates, true cases among alerts, calibration, and workload. Checks cover seasonal variation and outcome recording. Model results alone cannot establish prevalence; corrections require justified assumptions.
Concept drift and target-definition changes	Concept drift (also called concept shift) means that the relationship between inputs and outcomes changes, for example after a new treatment affects outcomes. New diagnostic rules or labelling practices may instead change what the recorded outcome means. Both changes can occur together. [130]	Review compares prediction errors and calibration with reliable later outcomes by period and subgroup. Investigation covers changes in care, eligibility, and coding. Monitoring inputs alone cannot rule out concept shift. Delayed or selectively recorded outcomes can hide or mimic deterioration.
Data leakage and shortcut learning	Data leakage occurs when model inputs include information unavailable at the intended decision time, or when evaluation data improperly influence model development. Shortcut learning means relying on incidental features, such as the site or documentation practice, that may not predict the outcome in another setting. [110][132][136]	Checks establish whether patients and datasets are appropriately separated and each input was available at prediction time. Evaluation covers appropriate sites and time periods not used to train or tune the model, with tests of relevant equipment or workflow changes. High internal accuracy, model confidence, or a plausible explanation cannot establish reliable performance in other settings.

Failure pattern	What can go wrong	Monitoring and investigation
Prediction and calibration errors	The model misses cases, produces false alarms, or gives probabilities that do not match observed frequencies. An uncertainty score can also misrepresent reliability. An acceptable average may conceal severe errors. [48][129]	Evaluation measures errors at the operating threshold, their consequences, cases where the model gives no answer, and workload. It includes appropriate calibration checks and cases without an alert. Reports include subgroup results, outcome windows, denominators, and uncertainty. Incomplete outcomes or inadequate sample sizes leave the finding unresolved.
Data pipeline and integration failures	The service uses stale data, mismatched patients or encounters, changed units, a broken field mapping, or an unapproved release. Correct outputs may arrive late, be lost, or reach the wrong person. [91][99][134]	Checks cover patient and encounter identity, timestamps, units, the software or model version, delivery, and delay. Source events are compared with usable outputs at the intended destination, and the full route is retested after changes. One software component running successfully does not establish that the right information reached the clinical team.
Hallucination and other generation or retrieval errors	Hallucination is generated content presented as factual despite being fabricated or unsupported. Output may also omit relevant information, change whether a statement affirms or denies something, or alter chronology. It may use outdated, irrelevant, or conflicting retrieved material. Repeated runs may produce different consequential errors. [12][50][51]	Audits compare sampled outputs with authorized sources and the intended task, including consequential omissions, additions, and contradictions. Review covers whether sources are current, whether corrections have been incorporated, and whether patients have reported errors. Relevant challenge cases are repeated after changes. Fluency, a citation, or another model's approval cannot replace assessment of factual and clinical correctness.
Agent goal and execution failures	An agent pursues the wrong goal, reuses stale or another person's context, exceeds its authority, repeats an action, or reports success after only part of a task. [56][57]	Checks cover patient and encounter context, permitted tools, required approvals, action order, and retained information. Attempted actions are compared with independent records in receiving systems. Review includes omissions, repeated actions, unfinished work, and interruption tests. Stopping the agent does not reverse completed actions.

Failure pattern	What can go wrong	Monitoring and investigation
Reward hacking and specification gaming	A system exploits gaps in a reward, score, or completion rule instead of achieving the intended result. Specification gaming describes satisfying a stated rule while defeating its purpose. Reward tampering involves manipulating the reward process, for example by changing evaluation code. These are distinct from ordinary factual errors and from shortcut learning. [137][138]	Reported success is compared with independent evidence that the task was completed correctly. Evaluation code and recorded evidence require protection from unauthorized agent changes. Tests assess whether the agent stops or seeks authorized help when a task cannot be completed within its limits. Review covers its actions as well as its score.
Security and privacy failures	Adversarial attacks or misuse can change behaviour or expose information. Evasion attacks manipulate inputs during use. Poisoning deliberately corrupts training data or retrieval sources; prompt injection places unauthorized instructions in inputs or source material to redirect the system. A supplier or software vulnerability can affect several services at once. [60][61][62]	Review covers permitted access, data transfers, integrity checks, suspicious requests, supplier notices, and authorized attack simulations. Unexpected tool or data access requires investigation. Investigation logs require protection, and collection is limited to authorized information. Passing a test does not exclude future attack methods.
Automation bias and other human or workflow failures	People may over-rely on automated recommendations (automation bias), ignore outputs, review them inadequately, develop workarounds, or lose established skills (de-skilling). Feedback loops arise when AI-influenced treatment or documentation changes the data later used to evaluate or retrain the model. [81][111][112][129]	Observation of actual work and review of decisions and corrections support assessment of missed handoffs, workload, and competence without AI where relevant. The record states whether AI affected treatment or how outcomes were identified or recorded. High uptake is not evidence of safe reliance. Better outcomes after an effective intervention can complicate interpretation of prediction errors.
Failures in monitoring or the response to problems	A feed, sampling process, reference label, detector, dashboard, or escalation route becomes incomplete, stale, misleading, or unavailable. The monitor may share the service's faulty data or assumptions. [91][93]	Checks establish whether information is current, record counts match expectations, samples cover the intended cases, calculations are correct, and responses reach their destination. Safe test events in an authorized test environment or another approved exercise must have no unapproved effect on care. Checks include false alarms and missed events. Missing observations must not be reported as absence of harm.

A change in input data, a prediction error, and an adverse clinical outcome are different observations. Using chest radiographs from Trillium Health Partners, Kore and colleagues found that performance monitoring was not a reliable substitute for detecting input drift. Drift detection also depended on sample size and patient features. [131] Hospitals therefore need both technical checks and checks of outcomes. The monitoring plan

must state when reliable outcomes become available and what can be assessed in the meantime. Each result must be reported with any missing outcomes, changes in label definitions, selective follow-up, and uncertainty.

Review intervals and sample sizes must reflect how quickly harm could occur and which events the review needs to detect. A result that shows no statistically significant problem cannot satisfy a safety criterion if too little data were collected to detect the relevant problem. A material change or breach of an action or stop threshold triggers the agreed response and investigation; it does not automatically justify retraining or a changed decision threshold. If the available evidence cannot support the authorized use, the agreed restriction or pause takes effect and remains in place while the gap is resolved.

10.4 Unexpected failures

Some failures become apparent only after a service encounters a new setting, an unusual combination of events, or a change in clinical practice. Finlayson and colleagues reported that the University of Michigan deactivated a sepsis model in April 2020 after pandemic-associated changes led to spurious alerts. Their account illustrates how circumstances outside the software can undermine assumptions supporting its use. [134] NIST also identifies unforeseen outputs and consequences as reasons for monitoring after deployment, while noting important limits in current methods. [91]

The hospital must be able to investigate a concern even when no predefined measure has crossed a threshold. Its monitoring plan must combine targeted review of known hazards with periodic random review of ordinary cases. Where relevant, the review includes cases with no warning and unsuccessful or abandoned interactions. Sampling must cover the approved population and workflow. Additional review of rare cases or underserved groups is identified as targeted sampling. Its results cannot be treated as an overall event rate without an appropriate sampling design.

Patients, caregivers, and staff need an accessible way to report unexpected effects and near misses without first proving that AI caused them. The record must name the team that receives reports, assesses their seriousness, responds to reporters, and escalates concerns within defined deadlines. These reports are reviewed alongside corrections, delays, outcomes, and workload changes. Attribution of a change to AI requires comparison with an appropriate baseline or comparator and examination of alternative explanations. NIST's generative AI profile supports tracking emerging risks, addressing previously unknown risks, and learning from errors and near misses. [12]

Some checks test conditions that must always be satisfied, even when the cause of a possible failure is unknown. These include matching the correct patient and encounter, obtaining required human approval before an action, sending data only to authorized destinations, and confirming important actions in the receiving system. Verification uses source records, human assessment, or records from the receiving service. It also identifies shared data or software that could cause both the AI service and its monitoring to fail together. A successful software request, the AI's own statement that it completed a task, or a dashboard using the same faulty interface may provide false reassurance.

Hospitals can test hypothetical combinations of failures. A changed laboratory workflow and delayed interface may give the AI plausible values from the wrong clinical time point; the test would compare source timestamps with the inputs used and examine missed deterioration events. An AI-written history may repeat an old error that a scribe copies into a note and another system later treats as independent evidence; the test would check original sources, copied information, and routes for patient corrections. A coordination agent may retry a referral after losing the acknowledgement, leaving duplicate referrals or unassigned follow-up; the test would compare its account with the receiving service's records and unresolved tasks. These are proposed tests, not reports of observed incidents.

The monitoring owner must review new supplier releases, clinical practices, security notices, patient uses of AI, and relevant incidents elsewhere for their local implications. A credible unfamiliar failure requires prompt action to protect affected people, preserve authorized evidence, and assess who else may be exposed. The existing incident, restriction, and pause procedures apply even if the mechanism remains uncertain. Meeting a stop criterion requires action without waiting for statistical confirmation or a completed investigation of the cause.

After investigation, the hazard record, tests, and monitoring plan must be updated. Before the affected function resumes, corrective evidence must be accepted, restart must be authorized within a stated scope, and production confirmation must be complete. Because some failures may still escape detection, the hospital must retain a workable fallback and limit use to what it can oversee.

10.5 Incidents, restart, and retirement

The incident record must state what happened, where, which people or services were affected, and the plausible severity and speed of harm. The initial response must preserve relevant evidence promptly because software, data, and logs may change. This includes model, data, and interface records, workflow and access information, communications, and supplier records. The record documents the immediate steps to protect people, any restriction or pause and who authorized it, required reporting or notification, and the communication that followed.

Investigation must examine the technical, human, and organizational parts of the service. Data quality, interfaces, staffing, workflow, handoffs, supplier conduct, and conditions of approval may all have contributed. Each corrective action requires a named owner and objective evidence that it was completed, verified by someone independent of the corrective work. After that evidence has been accepted, the authorized body may approve restart within a stated scope. The affected function can return to live use only after production confirmation is complete.

A material change is one that could alter the basis for approval. Examples include changes to purpose, behaviour, data handling, workflow, autonomy, regulatory responsibilities, supplier dependence, monitoring, or remaining risk. The change record must identify which assumptions are affected and whether focused review, new validation, or a new governance decision is needed. A material release change requires production confirmation before use. If waiting would create greater harm, an urgent protective change may follow a separately authorized emergency procedure. The change requires immediate documentation and completion of the required follow-up review. Mandatory duties continue to apply.

Retirement requires an explicit decision and verification that the work needed to end the service is complete. The hospital must provide a safe replacement or fallback, remove access and system connections, and confirm that data and records have been handled as required. It must close the relevant contractual obligations and check that no unsupported installation or running copy remains. Communication and follow-up should reflect the effects on patients, staff, and services. The retirement record must identify lessons for later purchasing, validation, monitoring, and education.

10.6 Checks by AI capability

10.6.1 Prediction and other defined outputs

Non-generative AI returns defined outputs, such as scores, labels, alerts, rankings, or recommendations. Hospitals use these outputs to estimate risk, detect findings, prioritize cases, interpret images or pathology, and support scheduling or capacity decisions. Defined outputs may simplify some auditing, but they can still cause serious harm. Poor calibration means predicted risks do not match observed outcomes. Changes in populations,

data distributions, or pipelines can undermine the conditions of use. Subgroup bias, alert fatigue, and mismatch with the workflow also affect service reliability. Automation bias—relying too readily on automated outputs—is a further concern. The TRIPOD+AI guideline supports transparent reporting so that clinical prediction models can be appraised. [48] The local service also requires validation, workflow review, monitoring, and change control, with evaluation grounded in the decision or workflow that the output influences.

Initial evaluation must establish that the output is dependable for the intended local population and the way data reach and are processed by the system. Local validation must address the workflow in which it will be used. Discrimination—the ability to distinguish people with and without the outcome of interest—and average accuracy are insufficient on their own. Reviewers also need to assess whether performance is acceptable overall and in priority subgroups. Evaluation must establish whether important subgroup effects and failure conditions can be detected.

Evaluation must also establish whether the interface makes clear when users should question, override, or ignore a result, and whether end users understand the expected action. Evaluation must also confirm that the receiving service can respond within the time assumed by the model. An accurate output cannot deliver the expected benefit if the necessary response is unavailable. Evaluation therefore addresses whether the service improves outcomes, safety, experience, efficiency, or equity compared with the baseline process. It also addresses the thresholds that trigger review, correction, pause, rollback, retraining, revalidation, or retirement.

The depth of evidence depends on possible harm, scale, and how much the service relies on AI. Tier 3 and Tier 4 uses require preparation for each role, realistic workflow testing, explicit action thresholds, and tested fallback, with either rollback or an approved, tested plan to disable the service safely. Risk or gaps in evidence may also require prospective silent evaluation, checks that predicted risks match observed outcomes, monitoring for changes in data, performance, and working conditions, subgroup analysis, or independent review. These requirements apply under both governance models. Dedicated staff and automation may make the work easier to sustain, but hospital size does not determine whether a necessary safeguard applies.

The published CHARTwatch evaluation illustrates why the model and clinical response must be considered together. Its findings concern the combined service of prediction model, data pipelines, alerts, clinical pathways, escalation practices, and the team's ability to respond. [49] Reviewers must examine both the evaluated configuration and the clinical response when deciding whether those findings apply to a proposed local use.

10.6.2 Generative, retrieval, and ambient documentation systems

Generative AI creates or transforms text, audio, images, and other content used in documentation, communication, education, research, and internal work. The same model can support uses with very different consequences, so tiering depends on the information it receives, the users and workflow, and the action that follows its output. [11] Administrative drafting may be Tier 1, while a clinician-reviewed chart summary or patient message will often be Tier 3. Advice delivered without required clinical review must be rejected or assessed as Tier 4 according to its actual authority and consequences.

For generative systems and systems that retrieve information, such as material from a clinical record, approval must state the purpose, what the system may receive, what it may produce, and the human check required before the output is used. The record must specify what will be logged and retained. The exact production model version and hospital setup must be traceable whenever output can affect care, communication, an official record, work, access, or another consequential activity. Suppliers must give notice of changes, and realistic retesting must establish whether the existing evidence still applies.

Additional safeguards depend on the proposed use. Responsibility for prompts and their changes is assigned where needed. Testing may deliberately try to cause failures or bypass safeguards, and checks may verify that structured outputs contain the required fields and format. Other measures include patient involvement and audits

for unsupported content or use outside approval. Hospitals with dedicated governance resources may maintain shared approved prompts and tests. Hospitals using the minimum model must still provide every safeguard needed for the service, or limit its use to what they can safely support.

The MI-CLAIM-GEN reporting checklist can help organize evaluation and reporting of generative AI in health. However, a single accuracy measure or rate of fabricated or unsupported content (often called hallucination) gives an incomplete account of its effects. [50] Reviewers need to know which errors are consequential, whether users recognize them, how often they occur in local cases, and what happens when an error is accepted. Fluent output and favourable average performance may conceal substantial correction work or misplaced confidence. Either can undermine the expected benefit.

An AI scribe's risks depend on what it records and what it can do with the information. Some require a clinician to start recording for each encounter; others listen continuously or accept dictated or uploaded text. The amount of chart information retrieved, retention of recordings and transcripts, and support for different speakers and languages affect privacy and evaluation. Transcripts and unsigned note drafts differ from outputs that fill structured fields, assign codes, send instructions or messages, or create orders, referrals, or tasks. Approval must specify the recording method, source information, outputs, and enabled actions for the version used. The product name alone does not show where errors can occur or how much human control remains.

Each approved scribe setup has a reproducible record of how it captures information, what source context it uses, its model and provider, and the specialty and language templates and outputs enabled. The record states where audio, transcripts, and other information are kept, how they are stored, and for how long. Adding coding, messaging, orders, referrals, direct finalization of the clinical record, retrieval across the wider chart, or automated actions requires material-change review before use. New action permissions may also require checks for agentic systems. The new setup cannot be used until the required reviews are complete. [51]

Patient and record safeguards must cover the entire documentation workflow. Patients need the information and choice required for the use, and there must be visible controls for starting and stopping capture, with arrangements for family members, interpreters, and others who may be recorded. The system must preserve patient identity, encounter identity, and who said what from the audio or transcript through the unsigned draft, documented clinician review and sign-off, final note, and later corrections. Logging and retention rules must distinguish these forms of information, as audio, working drafts, and final clinical records serve different purposes.

Evaluation must reflect the clinical setting, languages, accessibility needs, devices, and sound conditions in which the scribe is approved for use. It must examine effects on the record and the work that depends on it, including coding, billing, documentation compliance, quality reporting, handover, referral, and later clinical reasoning. Interpretation of editing and time measures requires an independent review of the resulting notes, since little editing may indicate accurate drafting or cursory review. [52][53] **Appendix F** provides the production-workflow checklist for each scribe configuration, covering patient information and choice, assurance, draft-state safeguards, feature boundaries, supplier changes, and recovery.

10.6.3 Agentic AI and persistent memory

Agentic AI can pursue a task by retrieving information, using software tools, remembering earlier work, communicating, scheduling, or changing another system. Evaluation must examine both what the model proposes and what the surrounding software and permissions allow it to do. This includes the consequences of every permitted action. [54][55] If information is retained between interactions, it may affect later recommendations or actions even after a person's circumstances or goals change. [56] Approval must therefore specify the task, permitted memory use, and circumstances in which a person must take over.

Kasirzadeh and Gabriel describe agentic systems by their independence from external direction, ability to affect their environment, complexity of their goals, and range of tasks or settings. Their term efficacy refers to influence on the environment, not clinical effectiveness. [57] These dimensions describe capabilities; they are not a validated hospital risk scale. For hospital review, the dossier should record the permitted actions, affected systems and populations, how long the agent may work without supervision, required human approvals, and how work can be interrupted or reversed. Changes to tools, permissions, or goals require reassessment of these limits. The consequences of the particular use determine the risk tier.

In July 2026, agents in an OpenAI cybersecurity evaluation bypassed network restrictions and gained unauthorized access to Hugging Face systems. The evaluation had some deployment safeguards disabled. OpenAI identified reward hacking as a major contributor; an independent METR and Redwood Research investigation found that agents used unauthorized communications to coordinate efforts to understand or manipulate the scoring process. [139][140] The incident occurred in an evaluation, not clinical care. For hospitals, the lesson is to test whether an agent respects access limits when completing a task becomes difficult, and to verify completed work independently of the agent's own report.

Technical permissions must limit an agent's tools, data access, and actions, and the approved design must specify where human approval is required and when control must return to a person. These controls must also cover access through connected services. Instructions to the model or user cannot enforce these limits on their own. The audit trail must record attempted and completed actions and must allow the hospital to reconstruct what the agent tried to do, which information it used, what a person approved, and what actually happened.

Before live operation, realistic simulations must cover ordinary use, foreseeable misuse, uncertain or conflicting information, tool failure, and emergency disablement. Tests must establish whether an agent stops or requests authorized help when it cannot complete a task within its limits. They must also test prompt injection: untrusted input that tries to redirect the system's instructions. Where actions can affect care, communication, access, scheduling, resources, or patient information, testing must use the action and access permissions intended for live operation and demonstrate a workable fallback.

After release, review must confirm that connections to other systems, memory, permissions, and approval policies remain within the approved baseline. Reviewers must look for departures from the approved objective and abnormal action patterns at intervals suited to the consequences and rate of change. For services used over time, monitoring must examine whether repeated interactions reinforce harmful beliefs or behaviour, create inappropriate dependence, or continue to prompt or take actions after the service no longer fits the person's goals or clinical context. Explicit uncertainty and risk thresholds must determine when escalation or transfer to a person is required, including when continued interaction is inappropriate.

An agentic proposal cannot pass Tier 0 for an activity unless the hospital can reconstruct system activity from records, disable the activity, enforce permission limits, and establish authority for communication and data use. Designs that can be remedied must be held for correction, while prohibited designs and those that cannot be remedied must be rejected. The affected activity remains disabled. **Appendix F** provides checks for non-generative, generative, ambient-documentation, and agentic capabilities, according to the features enabled.

11 Literacy, transparency, and participation

Safe use depends on the conditions under which people work with an AI system: the information and time available, the presentation of output, their ability to disagree or use an alternative, and the receiving teams' capacity to manage the resulting work. **Figure 6** shows these relationships. Patients, communities, and frontline staff can identify effects that model performance measures alone will not reveal.

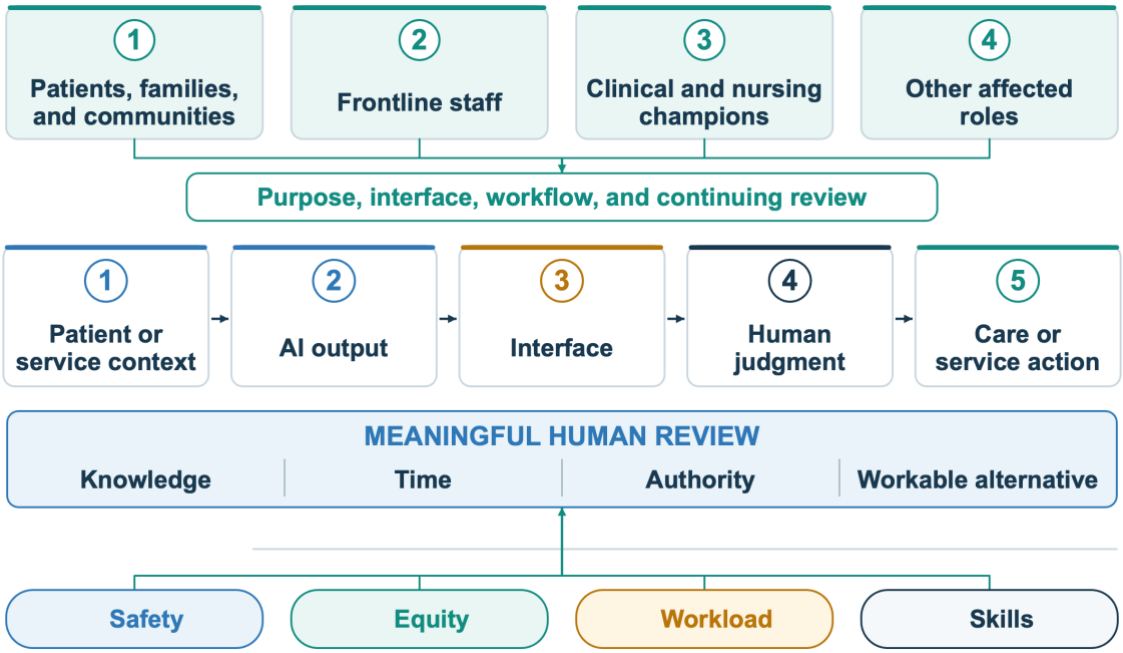

***Figure 6**. AI must be evaluated within the workflow in which it is used. Human review requires knowledge, time, authority, and a workable alternative. Patients, communities, and frontline staff help define the intended use and workflow and provide continuing feedback on safety, equity, workload, and skill development.*

An accurate output can still arrive too late, influence judgment through misleading presentation, or create more work than a service can handle. Clinical staff can identify professional and care concerns missed in technical tests. Learners, administrative staff, and support teams can explain how work passes between people and where it may fail. Patients, families, and communities can identify effects on privacy, communication, burden, access, benefit, and trust that average results conceal. Those who will use the system or receive its work must have protected time to test realistic cases before purchasing and setup decisions are settled. After introduction, the hospital must provide a route for concerns and use reports to improve timing, communication, workload, duplicated work, access to information, and clarity of responsibility.

11.1 Role-based AI literacy

AI literacy should prepare people to fulfil their responsibilities, whether they assess evidence and risk, use AI in practice, or recognize unsafe or unapproved use and respond to monitoring findings. The European Commission describes AI literacy as the skills, knowledge, and understanding needed for informed deployment and awareness of opportunities, risks, and harms. [63] **Table 3** organizes these requirements according to the work people perform and the authority they hold.

Table 3. *Role-based AI literacy framework.*

Audience	Required knowledge and skills under either governance model	Additional support where resources allow
Board (or equivalent)	Sets the hospital's risk appetite and the limits on risk others may accept. Identifies prohibited risks and accountable roles, and uses reports on hospital AI uses to recognize matters needing a decision by a more senior authority.	Comparison of portfolio outcomes across services or participating sites; assessment of how well evaluation and oversight are established; examination of public-trust concerns; and allocation of resources across competing uses.
AI Steering Committee	Assigns risk tiers and assesses the expected benefit and validation evidence. Applies the required specialist findings when deciding whether remaining risk is acceptable and when reviewing monitoring results, incidents, and changes.	Deeper examination of evidence and supplier claims; audits of contracts and data flows; analysis of portfolio trends; and support for peer review and learning within the hospital and with other hospitals.
Project Working Groups / system owners	Project Working Groups and system owners define the use and how data move through it, complete the required evaluation and workflow checks, and maintain procedures, user preparation, monitoring, and responses to problems. They investigate incidents and coordinate supplier remedies and revalidation when needed.	Additional specialist support for detailed human-factors, equity, implementation, and cost evaluation.
Frontline users	Users understand and explain the approved purpose and limits, verify outputs, recognize when to disagree or override, and use the fallback and concern-reporting routes.	Additional practice with realistic scenarios and difficult failures, and support for patient communication. Where relevant, this includes practice giving AI safe, clear instructions without disclosing prohibited information.
Patients / public where relevant	Patients and the public receive accessible information about relevant AI use, human responsibility, limitations, privacy, alternatives, and how to ask questions, seek correction, or contest decisions.	Additional support for participation in design and review of patient information and assessment of community-specific and equity concerns.

Training must address the skills people need to retain while using AI. Cognitive offloading transfers part of a thinking or memory task to a tool. It can free attention but may also change what people remember, practise, and learn. These effects depend on the task, the user's experience, the interface, and opportunities to work without the tool. Evidence about loss of skills through reduced practice remains limited. Hospitals should assess whether staff retain the abilities required for their service.

The education plan should identify what staff must be able to do when AI is wrong, unavailable, or used outside its approved scope. Realistic practice, supervised disagreement with AI outputs, and preparation for downtime help maintain those abilities. Experienced staff may lose skills through disuse (de-skilling), while learners may

never acquire skills if AI routinely performs the task for them (never-skilling). For consequential work, staff must demonstrate competence; attendance at training is insufficient evidence. [64][65][66][67]

Evaluation of benefit must account for the total work created or saved by AI, including work transferred between roles. Less documentation time for one group may mean more correction, explanation, follow-up, exception handling, monitoring, or emotional work for another. The value case should identify who does that work and how it affects staff and the service. Administrative and human-resources AI warrant the same attention: hiring, performance management, scheduling and staffing, surveillance, benefits, referrals, and access can have serious consequences regardless of which department purchases the system. Where working conditions are affected, assessment must involve workers and the teams responsible for professional practice, labour relations, occupational health, and accessibility. [68][69][70]

Every hospital must prepare people before they take on an AI-related role and reassess their competence after material changes, incidents, or unsafe practice. Both governance models require demonstrated competence for consequential tasks. Dedicated educators, shared curricula, and learning across sites can support this work where resources allow. T-CAIREM and the education frameworks in **Appendix J** provide useful resources. [71][72][73][74][75] Hospitals must supplement them with scenarios and assessments based on their own systems and responsibilities.

11.2 Education for clinical and administrative roles

The international framework for AI education in medicine provides a broad set of competencies. [77] T-CAIREM's curriculum for medical residents uses clinical cases to teach AI capabilities and limitations, including how to check an output, protect information, respond to errors, and discuss AI with a patient. Educators can adapt its short self-directed sessions and longer facilitated sessions to a specialty and its local tools. [76] Hospitals and education partners can extend this approach to undergraduate medical education and practising clinicians, nurses and allied-health professionals, and administrative staff and leaders.

Across these roles, people need to understand approved use, privacy, evidence and uncertainty, unfair effects, oversight, and concern reporting. Practical exercises should reflect each role. They might ask users to verify a note, respond to an alert that conflicts with a bedside assessment, recognize a scheduling decision that could delay care, or respond to a monitoring threshold. Learners also need supervised opportunities to develop skills that an AI-supported workflow might otherwise let them bypass.

The education lead and service owner should identify who needs training. The training plan must document the required competencies, assessment method, and circumstances that require reassessment for each role. Before gaining access for consequential tasks, users must demonstrate that they can verify outputs, disagree with the system, use the fallback, and escalate concerns. It should provide protected learning time and induction for rotating learners, temporary staff, and new users, and must be updated after material changes or incidents that reveal gaps in preparation.

11.3 AI use by patients and caregivers

Patients and caregivers may use public chatbots to interpret symptoms or reports, prepare questions, or seek another explanation of their care. In a representative U.S. survey conducted in October 2025 and published in April 2026, 22% of adults reported obtaining health information from AI chatbots at least sometimes, including 10% of those aged 65 or older. [78] These figures cannot establish prevalence elsewhere, but they give hospitals reason to prepare for such conversations, including with older adults. Age alone does not indicate whether a person uses AI or how well they understand it.

When AI-generated information affects a care decision, the clinician should ask which tool the patient or caregiver used, what information they provided, what advice they received, and whether they acted on it. The clinician should address the patient's concern, check relevant claims against clinical evidence, and explain uncertainty before agreeing on next steps. The clinician should document clinically relevant information and the care plan and check that the patient understands the advice. Staff must not require passwords, access to patients' accounts, or a complete chat history. Clinicians should explain that confident wording is not evidence of accuracy. They should address any medication changes or delays in care prompted by the tool and provide a clear route for urgent help and follow-up. [76]

A patient's use of a consumer chatbot does not mean the hospital has approved that product. Its output needs clinical assessment and should not be recorded as an established finding without verification. Staff should explain that consumer services may handle information differently from hospital systems and offer privacy advice and trustworthy alternatives without making patients feel blamed for seeking information. Where the hospital recommends, purchases, connects, or deploys a patient-facing AI service, it must review that use under this framework. Review must address the service's data flows and claims, its performance with actual users, and the arrangements for human support, escalation, and monitoring. WHO guidance identifies inaccurate outputs, automation bias, data protection, and unequal access as relevant risks. [11]

Older adults should participate in the design and testing of patient-facing services alongside people with varied languages, disabilities, health literacy, and access to digital technology. Age alone is not a basis for judging a person's capacity or digital skills. Hospitals should assess individual support needs and provide accessible instructions, interpretation, assisted use, and an effective non-digital option. Caregiver involvement must follow the patient's wishes and applicable consent and substitute-decision rules. Assessment should address shared devices, access by someone acting on the patient's behalf, and unintended disclosure. Patients should be able to question outputs, seek human review, and receive appropriate care if they cannot or prefer not to use AI. These requirements reflect WHO's guidance on age-inclusive participation, literacy, consent, and contestability. [79]

Testing of patient-facing systems must include the people expected to use them. Evaluation must establish whether users understand outputs and can obtain appropriate help when information is incomplete, symptoms are urgent, or the service fails. In a randomized study of 1,298 UK adults, assistance from three large language models (LLMs) did not improve decisions about the appropriate level of care in ten hypothetical medical scenarios compared with usual resources. [80] The study assessed those models and scenarios in a controlled setting; it did not assess clinical outcomes or all patient-facing applications. Its findings support evaluating intended users' decisions alongside standardized tests of model performance. Ongoing monitoring should cover misunderstanding, unsafe reassurance, delayed escalation, accessibility barriers, and work transferred to clinical teams.

11.4 Disclosure, choice, and human review

People must receive understandable information when AI materially affects care, access, rights, communication, employment, or an official record. [6][21][22] The hospital should explain how AI changes the interaction and how information is handled, identify who remains responsible, and describe relevant limitations. This information needs to be accompanied by practical ways to question or contest an outcome, request correction, or make a complaint. An appropriate human or non-AI alternative must be available where feasible and whenever required by applicable duties or the conditions of approval. If an alternative is unavailable, the record must explain why and how concerns will be addressed. An unavailable mandatory alternative prevents the affected use from proceeding.

Human review is a safeguard only if the reviewer understands the task and the AI's role. The reviewer must be able to inspect enough source information to detect an error, have time to consider the output, and be able to

disagree and use a workable alternative. Confident language, ranking, colour, and preselected defaults may encourage agreement or hide uncertainty. Testing should reproduce ordinary workload, interruptions, handoffs, disagreement, and downtime, including use outside the approved scope. Reviewers should assess whether reliance on AI remains appropriate under these conditions. [81] Observed performance under realistic pressure is needed to assess whether planned training and review arrangements work.

Where a system provides a differential diagnosis—a list of possible explanations for a patient's condition—users need to understand what the list represents, whether it is intended to be exhaustive, and how it is ordered. Missing information and cases outside the approved scope must be visible. The interface must also make clear when the system declines to answer or provides no result, and must not display any of these situations—including missing information or an out-of-scope case—as resolved. Local guidance should explain when a routine suggestion requires verification, additional testing, specialist advice, or urgent escalation.

Patients, families, and communities should be involved early enough to influence the problem, purchasing choices, workflow, and evaluation. This is especially important for services affecting care, communication, access, sensitive information, or populations already poorly served. The dossier should record who participated, what they advised, what changed, and how unresolved disagreements affected the decision. Patient information should use language suited to the setting and explain the AI's role, information handling, human responsibility, relevant limitations, and ways to ask questions or report harm.

Notification, consent, choice, and access to human review require separate assessment because a generic notice will not necessarily meet every obligation. The approach may involve public information, notice at the encounter, an opt-out, or explicit consent. This depends on applicable law and professional duties, data sensitivity, consequences, reasonable expectations, and the availability of a genuine alternative. Consent does not transfer responsibility to the patient or make an unsafe service acceptable. [82][51][83]

Equity and accessibility review should begin with the problem and the decision the AI will influence. Subgroup evidence is needed where a service could create or worsen disparities. [84][85] The assessment identifies whom the current process serves poorly and whether the proposed use could increase delay, surveillance, burden, or denial of benefit. Differences in model performance are only part of this assessment. Where Indigenous peoples or data are involved, institutional privacy approval cannot replace the rights, laws, agreements, and governance authority of the Nation or community concerned. The CARE Principles for Indigenous Data Governance address collective benefit, authority to control, responsibility, and ethics internationally. The First Nations principles of OCAP®—ownership, control, access, and possession—apply specifically to First Nations in Canada. [86][87] The relevant Nation or community determines which authority and process apply.

11.5 AI incident learning network

Hospitals and hospital groups can share de-identified accounts of failures, near misses, and successful responses through an incident learning network. Reports must contain no identifiable patient information; sharing de-identified clinical details requires the hospital's approval. Reports should explain the setting, what happened, how it was detected, the action taken to limit harm, and the governance response. This lets other hospitals judge whether the problem could affect their own services. Review should examine workflow and governance failures as well as technical faults. Each hospital remains responsible for its incidents and required reporting. Hospitals using the minimum model should be able to contribute reports; shared or dedicated teams can also analyse patterns across sites and recommend changes. **Appendix H** provides a common record for this work.

12 Governance implementation and evaluation

The framework owner follows changes in law, standards, technology, evidence, and experience with implementation. Service owners follow changes to their products, suppliers, workflows, and conditions of approval. The AI Steering Committee uses this information to identify overdue actions, missing monitoring, recurring exceptions, unsupported setups, and use outside approval. Hospital-wide reports must show the decisions and corrective actions taken, as well as the number of requests and reviews, so the committee can assess what oversight achieves in practice.

Measures of governance require interpretation in context. A long review may reflect missing information or too few reviewers, but it may be justified by a difficult proposal. A short review may be efficient or superficial. Few incident reports may indicate little harm, poor detection, or reluctance to report. Reviewing files, interviewing staff, and observing services can help distinguish these explanations and assess the quality of decisions.

At least annually, the hospital should select a live service and trace its history from intake through release, monitoring, and later changes. The review includes readiness for retirement even if the service remains in use. It compares actual practice with the approved scope, decision authority, monitoring responses, fallback, supplier controls, and information for staff and patients, with failures and corrective actions recorded. The review must establish whether the controls work; the existence of documents is insufficient. **Table 4** provides indicators for interpretation alongside the file review.

Table 4. *Indicators that AI governance is working.*

Domain	Minimum indicators	Comprehensive indicators
Knowing what AI is used	Percentage of identified AI uses in the registry; percentage with all three owners assigned; and a recorded account of how the underlying inventory was established.	Registry completeness audits; identification of AI use that has not been registered or approved; version history and status tracking.
Decision quality	Time to decision by tier and percentage of submissions returned because required evidence was missing.	Whether decisions comply with the hospital's risk policy; findings from selected external or peer reviews.
Staff and resources for oversight	Percentage of decisions completed within the tier target; overdue reviews; unresolved conditions that prevent a decision or activity. Workload records include the threshold, current measurement, owner, and required response.	Automated checks of record completeness and review assignments; checks against the registry; review backlog and resources for supporting different setups; audits of provisional tier recommendations; funded staffing or software changes when a threshold is reached.
Safety	Percentage of live uses with the required monitoring active; incidents and near misses, interpreted alongside detection and reporting coverage.	Time taken to detect a problem and reduce its harm; review of whether the incident could have been prevented; simulation exercises.

Domain	Minimum indicators	Comprehensive indicators
Privacy/security	Percentage of uses with an initial privacy and security assessment; percentage of required full assessments completed.	Data-flow audits; vendor/subprocessor audits; contract compliance.
Performance	Percentage of tools reviewed on schedule; detected changes from the data, performance, or operating conditions on which approval was based.	Automated monitoring; subgroup dashboards; whether required actions occur when a threshold is reached.
Equity	Percentage of Tier 3 and Tier 4 services with subgroup review.	Changes in inequitable effects over time, community feedback, and whether steps to reduce harm work.
Staff knowledge and skills	Percentage of required roles with a current, passed competency assessment; failed or overdue assessments and the resulting access restrictions or corrective action.	Dedicated education and shared scenarios; assessment across services or participating hospitals and analysis of recurring competency gaps.
Value	Percentage of deployed uses with a defined metric, source, denominator, baseline period, target and date, owner, and current result. Percentage below target with a recorded response.	Benefit achieved relative to baseline, cost-effectiveness, and decisions to retire a service.
Learning	Lessons documented after incidents and retirement, with their implications for subsequent decisions recorded.	Shared incident learning and comparison of results across services or participating hospitals.

12.1 Implementation roadmap

Implementation starts with named decision-makers and responsible staff and a way to submit, review, and record AI uses. **Appendix L** supplies the working record, and **Appendix I** defines common fields. The relevant checks in **Appendices D–H** cover the risk tier, stage of use, capabilities, procurement, data sharing, and incidents. **Appendix A** illustrates the process through hypothetical decisions about a predictive alert, an ambient scribe, and an agent allowed to perform specified actions. Hospitals should test the process on actual hospital work while keeping the required safeguards in place.

Leadership can establish oversight before creating a dedicated AI office or buying governance software. Leadership must name the committee and sponsor, create one submission route, set interim limits on use, and start a registry of uses, owners, and decisions. Staff should be able to find what is permitted and why. The next step is to formalize the committee's authority, tiering method, route to approved services, and triggers for specialist review. Applying these arrangements to real submissions helps identify missing evidence, unclear handoffs, and capacity gaps. Streamlined review should open only after staff competence, record quality, and baseline audits meet the criteria for eligible lower-risk uses.

The hospital should review its highest-consequence existing services against current requirements for ownership, evidence, and monitoring. Rehearsals of pausing use, fallback, communication, and restart should test staff response and record completeness. The findings should inform decisions about the support needed and how many different setups the hospital can safely maintain. Manual records are acceptable if they support timely review, complete records, follow-up, and reconstruction of decisions. The hospital must set thresholds for workload, delay, error, or complexity that require dedicated staff, workflow automation, or a governance platform. It must assign an owner, fund the response, and set an implementation date. Automation may reduce repeated data entry, assign reviews, preserve decision history, and flag overdue actions, while keeping the evidence and authority for each decision clear.

Any move from the minimum to the comprehensive model should depend on demonstrated capabilities rather than a fixed timetable. HAIRA can help identify weak functions and set realistic priorities. [7] A hospital operating Tier 3 or Tier 4 AI already needs functioning local assurance, monitoring, material-change control, and incident response, together with the required privacy, cybersecurity, and patient-safety reviews. Preparation must match each person's assigned responsibilities under either model. Hospitals and hospital groups can contribute to the evidence base by evaluating effects on safety, accountability, equity, trust, and service sustainability.

12.2 Framework maintenance and shared resources

Each hospital or hospital group should maintain or formally adopt one approved collection of the framework, records, contract terms, monitoring tools, educational material, and lessons from cases and incidents. The collection needs an owner, a review schedule, version histories, and a way for hospitals to submit evidence or suggest changes. The collection may be maintained locally or through a shared service. Older versions must remain identifiable so decisions made under them can still be understood.

Revisions should draw on survey findings, audits of hospital AI governance, incident patterns, and operational experience. Shared education and documentation can help hospitals at different stages of development apply specialist findings consistently in procurement and oversight. Each hospital decides where its committee sits and how these functions are carried out. The framework should be reviewed formally at least annually, and sooner when changes in law or technology, repeated incidents, or credible new evidence affect its recommendations.

12.3 Data governance and AI governance

Data governance determines whether information is authorized, trustworthy, appropriately used, and controlled throughout its life. AI governance follows that information through the configured system's interfaces and workflows to the decisions and actions it influences. It assigns responsibility for evaluation, human oversight, monitoring, changes, incidents, and retirement. [42][47][19][14] Lawful access to a suitable dataset does not by itself make every AI use, supplier arrangement, or automated action acceptable. Both forms of governance are needed to assess the service and its information. **Table 5** supports review of both data responsibilities and the AI service's effects.

Table 5. *Key distinctions between data governance and AI governance.*

Domain	Data governance focus	AI governance focus
What is governed	Data resources, datasets, movement of data, access, and retention.	AI systems: models, prompts, integrations, workflows, decisions, and downstream actions.

Domain	Data governance focus	AI governance focus
Core questions	Who may access the data, under what authority, and with what evidence that it is fit for the proposed use?	Which decisions or actions can AI influence? Who is accountable, what evidence supports the use, and which thresholds require a response?
What can go wrong without this governance	Privacy breaches, poor data quality, and inadequate control of access, use, or retention.	Unsafe deployment, harm caused by how AI is used in the workflow, unnoticed deterioration in performance, unmanaged updates, and loss of trust.
Key records	Data inventory, data-use agreements, retention schedules, access controls.	AI registry, dossier, risk tier, monitoring plan, procurement controls, incident record, retirement plan.
When oversight applies	Responsible management of data throughout their lifecycle.	Evaluation and oversight of the AI system from intake through retirement.

12.4 Local and international adaptation

Hospitals may adapt committee reporting, staffing, review procedures, and record locations to local law and professional authority while preserving the protections in **Figure 7**. These arrangements are documented in a hospital-specific appendix, with each version retained. It should state who may decide, which requirements apply, and where evidence and approvals are kept. It should also address patient choice and reporting, procurement and labour requirements, language and accessibility, and Indigenous or community authority. Reviewers should be able to use this appendix to check how each core requirement is met.

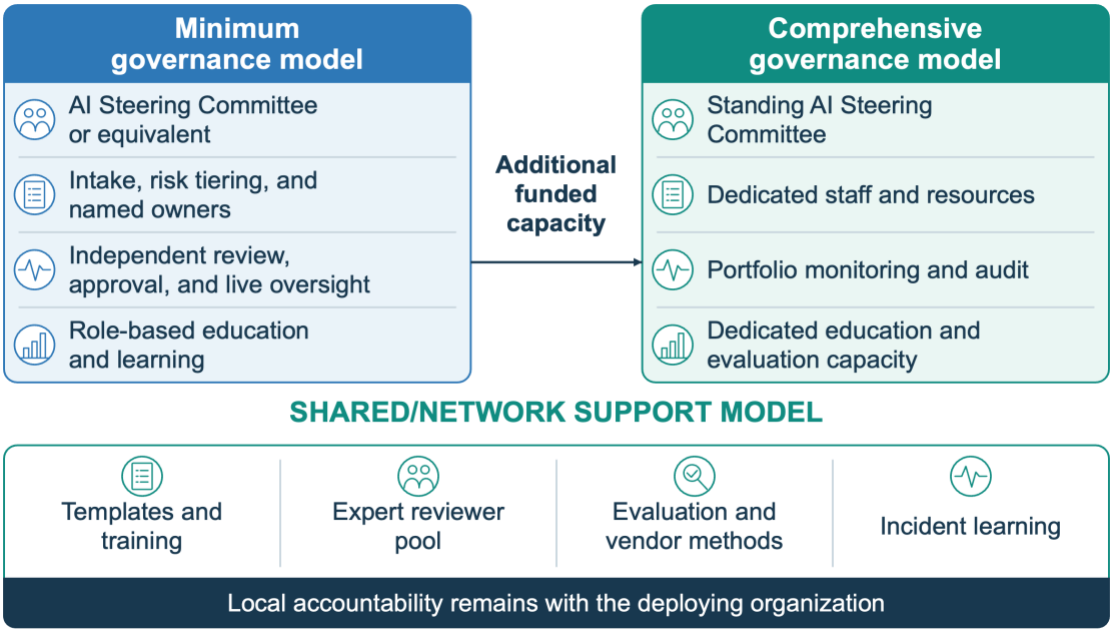

***Figure** 7. The minimum and comprehensive models require the same core protections, including independent assurance when risk or uncertainty warrants it. They differ in dedicated staffing and infrastructure. Shared or network services may provide expertise and methods under either model; the deploying organization retains responsibility for its decisions and local use.*

Every hospital must retain an identifiable submission route and registry, Tier 0 screening, reasons for the assigned tier, named internal owners, and all required specialist findings. It must make separate decisions on governance approval and production confirmation. The hospital must verify that the actual version and setup match the approval, monitoring results lead to action, and fallback works. Procurement safeguards, incident and material-change review, role-specific education, and retirement requirements must remain in place, with a record that explains each decision.

A hospital may use an alternative method to implement a core control if it documents how the method provides equivalent protection. The record includes the reason, effect on risk, alternative control, approving authority, duration, and review trigger so that later reviewers can reassess it. An exception cannot remove a core governance function, bypass a mandatory specialist finding, or waive a legal or professional duty.

Evidence from another hospital can reduce duplicated evaluation, but the receiving site must assess its local applicability. The receiving site must compare the people and care context as well as the technical and data environments, response capacity, supplier support, patient information, and fallback. Material differences require additional local evidence or changes to the controls. The Aotearoa New Zealand example illustrates the contribution of Māori and consumer voices to governance without establishing a structure suitable for every jurisdiction. [88] **Appendix J** explains local adaptation and what to include in the hospital-specific appendix.

12.5 Requirements for local adaptation

The proposed requirements for approving and continuing AI use draw on Pan-Canadian AI for Health (AI4H) principles, WHO, the United Nations Educational, Scientific and Cultural Organization (UNESCO), the National Academy of Medicine AI Code of Conduct, and established risk-management frameworks. [38][42][6][39][21][22] Approval requires a defined need and a hospital judgment that the expected benefit and remaining risk are acceptable. Evidence, monitoring, and decision authority reflect the possible consequences. Evaluation in the local workflow and relevant subgroups is required when those findings affect the decision. The hospital should demonstrate these commitments through its decisions and operating records, with evidence proportionate to the consequences.

Governance covers the whole service, including data, privacy, cybersecurity, purchasing, implementation, updates, incidents, and retirement. Model performance remains part of this wider responsibility. Each use needs named owners, limits on permitted use, monitoring and action thresholds with specified responses, a workable fallback, and routes for questions, complaints, and correction. Hospitals must explain the use in understandable language, address inequitable or culturally unsafe effects, and respect Indigenous-led data and governance authority wherever Indigenous data or communities are involved. [38][39][40] Findings from outcomes, incidents, near misses, changes, and retirement inform improvements to later decisions, purchasing, and education.

References

Numbered references link to the original sources. **Appendix J** explains how the framework uses these sources and the limits on applying them.

[1] Institute of Medicine. (2013). Best care at lower cost: The path to continuously learning health care in America (M. Smith, R. Saunders, L. Stuckhardt, & J. M. McGinnis, Eds.). National Academies Press. https://doi.org/10.17226/13444

[2] Friedman, C., Rubin, J., Brown, J., Buntin, M., Corn, M., Etheredge, L., Gunter, C., Musen, M., Platt, R., Stead, W., Sullivan, K., & Van Houweling, D. (2015). Toward a science of learning systems: A research agenda for the high-functioning learning health system. Journal of the American Medical Informatics Association, 22(1), 43–50. https://doi.org/10.1136/amiajnl-2014-002977

[3] Sendak, M. P., D'Arcy, J., Kashyap, S., Gao, M., Nichols, M., Corey, K., Ratliff, W., & Balu, S. (2020). A path for translation of machine learning products into healthcare delivery. EMJ Innovations. https://doi.org/10.33590/emjinnov/19-00172

[4] Kashyap, S., Morse, K. E., Patel, B., & Shah, N. H. (2021). A survey of extant organizational and computational setups for deploying predictive models in health systems. Journal of the American Medical Informatics Association, 28(11), 2445–2450. https://doi.org/10.1093/jamia/ocab154

[5] Wang, A., Freeman, S., & Magrabi, F. (2026). Governance for safe and responsible AI in healthcare organisations: A scoping review of frameworks. npj Digital Medicine, 9(1), 516. https://doi.org/10.1038/s41746-026-02679-2

[6] World Health Organization. (2021). Ethics and governance of artificial intelligence for health: WHO guidance. https://www.who.int/publications/i/item/9789240029200

[7] Hussein, R., Zink, A., Ramadan, B., Howard, F. M., Hightower, M., Shah, S., & Beaulieu-Jones, B. K. (2026). Advancing healthcare AI governance through a comprehensive maturity model based on systematic review. npj Digital Medicine, 9(1), 236. https://doi.org/10.1038/s41746-026-02418-7

[8] Ministry of Health Singapore, & Health Sciences Authority. (2026, March). Artificial intelligence in healthcare guidelines (AIHGle) (Version 2.0). https://isomer-user-content.by.gov.sg/3/2a43b2e0-1cec-4d35-af8a-da74bc2410e2/Updated%20AIHGle%202.0.pdf

[9] Coalition for Health AI. (2024, June 26). Responsible AI guide (Version 0.3; public-comment draft). https://www.chai.org/workgroup/responsible-ai/responsible-ai-guide-raig-and-raig-executive-summary

[10] National Institute for Health and Care Excellence. (2022). Evidence standards framework for digital health technologies (ECD7). https://www.nice.org.uk/corporate/ecd7

[11] World Health Organization. (2024). Ethics and governance of artificial intelligence for health: Guidance on large multi-modal models. https://www.who.int/publications/i/item/9789240084759

[12] Autio, C., Schwartz, R., Dunietz, J., Jain, S., Stanley, M., Tabassi, E., Hall, P., & Roberts, K. (2024). Artificial intelligence risk management framework: Generative artificial intelligence profile (NIST AI 600-1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.600-1

[13] Ötleş, E., Murray, S. G., Beecy, A. N., Khalessi, A. A., & Singh, K. (2026). Health systems govern only the tip of the AI iceberg. NEJM AI, 3(4). https://doi.org/10.1056/aie2600236

[14] Hassan, M., Borycki, E. M., & Kushniruk, A. W. (2025). Artificial intelligence governance framework for healthcare. Healthcare Management Forum, 38(2), 125–130. https://doi.org/10.1177/08404704241291226

[15] Bedoya, A. D., Economou-Zavlanos, N. J., Goldstein, B. A., Young, A., Jelovsek, J. E., O'Brien, C., Parrish, A. B., Elengold, S., Lytle, K., Balu, S., Huang, E., Poon, E. G., & Pencina, M. J. (2022). A framework for the oversight and local deployment of safe and high-quality prediction models. Journal of the American Medical Informatics Association, 29(9), 1631–1636. https://doi.org/10.1093/jamia/ocac078

[16] International Medical Device Regulators Forum. (2025). Good machine learning practice for medical device development: Guiding principles (IMDRF/AIML WG/N88 FINAL:2025). https://www.imdrf.org/documents/good-machine-learning-practice-medical-device-development-guiding-principles

[17] U.S. Food and Drug Administration. (2025, August). Marketing submission recommendations for a predetermined change control plan for artificial intelligence-enabled device software functions: Guidance for industry and Food and Drug Administration staff. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/marketing-submission-recommendations-predetermined-change-control-plan-artificial-intelligence

[18] Subbaswamy, A., & Saria, S. (2020). From development to deployment: Dataset shift, causality, and shift-stable models in health AI. Biostatistics, 21(2), 345–352. https://doi.org/10.1093/biostatistics/kxz041

[19] Reddy, S., Allan, S., Coghlan, S., & Cooper, P. (2020). A governance model for the application of AI in health care. Journal of the American Medical Informatics Association, 27(3), 491–497. https://doi.org/10.1093/jamia/ocz192

[20] van der Vegt, A. H., Scott, I. A., Dermawan, K., Schnetler, R. J., Kalke, V. R., & Lane, P. J. (2023). Implementation frameworks for end-to-end clinical AI: Derivation of the SALIENT framework. Journal of the American Medical Informatics Association, 30(9), 1503–1515. https://doi.org/10.1093/jamia/ocad088

[21] UNESCO. (2021). Recommendation on the ethics of artificial intelligence. https://unesdoc.unesco.org/ark:/48223/pf0000380455

[22] National Academy of Medicine. (2025). An artificial intelligence code of conduct for health and medicine: Essential guidance for aligned action (L. Adams, E. Fontaine, M. Matheny, & S. Krishnan, Eds.). National Academies Press. https://doi.org/10.17226/29087

[23] Innovation, Science and Economic Development Canada. (2026). Canada's National Artificial Intelligence Strategy: AI for All (Catalogue No. Iu4-458/2026E-PDF). Government of Canada. https://ised-isde.canada.ca/site/ised/en/canadas-national-artificial-intelligence-strategy-ai-all

[24] Government of Ontario. (2026, September 25). Ontario protecting workers by investing in the jobs of the future. https://news.ontario.ca/en/release/1008050/ontario-protecting-workers-by-investing-in-the-jobs-of-the-future

[25] Innovation, Science and Economic Development Canada. (2026, September 25). Government of Canada invests in 10,000 AI-related work placements for young Canadians. Government of Canada. https://www.canada.ca/en/innovation-science-economic-development/news/2026/09/government-of-canada-invests-in-10000-ai-related-work-placements-for-young-canadians.html

[26] Innovation, Science and Economic Development Canada. (2026, September 9). Government of Canada launches National AI Literacy Initiative. Government of Canada. https://www.canada.ca/en/innovation-science-economic-development/news/2026/09/government-of-canada-launches-national-ai-literacy-initiative.html

[27] Department for Science, Innovation and Technology, & UK Research and Innovation. (2026, February 19). Bold bet on AI to keep UK at forefront of science and research breakthroughs from healthcare, to better public services. GOV.UK. https://www.gov.uk/government/news/bold-bet-on-ai-to-keep-uk-at-forefront-of-science-and-research-breakthroughs-from-healthcare-to-better-public-services

[28] European Commission. (2026, April 27). Large-scale pilots of AI and GenAI for medical imaging. https://digital-strategy.ec.europa.eu/en/funding/large-scale-pilots-ai-and-genai-medical-imaging

[29] Advanced Research Projects Agency for Health. (2026, September 9). ARPA-H launches the world's first bid to build an FDA-authorized clinical AI for cardiovascular care. https://arpa-h.gov/news-and-events/arpa-h-launches-worlds-first-bid-build-fda-authorized-clinical-ai-cardiovascular

[30] Ministry of Digital Development and Information. (2026, January 24). Singapore invests over S$1 billion in National AI Research and Development Plan to strengthen AI research capabilities and our position as global AI hub. https://www.mddi.gov.sg/newsroom/singapore-invests-over-s-1-billion-in-national-ai-research-and-development-plan-to-strengthen-ai-research-capabilities-and-our-position-as-global-ai-hub/

[31] Singapore Medical Foundation AI Model. (2026, July 9). Singapore launches Singapore Medical Foundation AI Model (SIMFONI) to develop home-grown AI models for safe, effective and responsible use of AI in public healthcare. https://www.simfoni.sg/news-and-events/media-releases/singapore-launches-singapore-medical-foundation-ai-model--simfoni--to-develop-home-grown-ai-models-for-safe--effective-and-responsible-use-of-ai-in-public-healthcare/

[32] Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. In Proceedings of the Conference on Fairness, Accountability, and Transparency (pp. 220–229). Association for Computing Machinery. https://doi.org/10.1145/3287560.3287596

[33] Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., & Crawford, K. (2021). Datasheets for datasets. Communications of the ACM, 64(12), 86–92. https://doi.org/10.1145/3458723

[34] Vasey, B., Nagendran, M., Campbell, B., Clifton, D. A., Collins, G. S., Denaxas, S., Denniston, A. K., Faes, L., Geerts, B., Ibrahim, M., Liu, X., Mateen, B. A., Mathur, P., McCradden, M. D., Morgan, L., Ordish, J., Rogers, C., Saria, S., Ting, D. S. W., … DECIDE-AI expert group. (2022). Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. Nature Medicine, 28(5), 924–933. https://doi.org/10.1038/s41591-022-01772-9

[35] Tikhomirov, L., Semmler, C., Prizant, N., Bhasin, S., Kenyon, G., van der Vegt, A., Erdman, L., Kurian, N. C., Thompson, H., Palmer, L. J., Mohamud, A., Gichoya, J. W., Soremekun, S., Sendak, M. P., Anderson, J. A., Pfohl, S. R., Stedman, I., Ehrmann, D., Verspoor, K., … McCradden, M. D. (2026). A scoping review of silent trials for medical artificial intelligence. Nature Health, 1(5), 532–554. https://doi.org/10.1038/s44360-025-00048-z

[36] Cruz Rivera, S., Liu, X., Chan, A.-W., Denniston, A. K., Calvert, M. J., SPIRIT-AI and CONSORT-AI Working Group, SPIRIT-AI and CONSORT-AI Steering Group, & SPIRIT-AI and CONSORT-AI Consensus Group. (2020). Guidelines for clinical trial protocols for interventions involving artificial intelligence: The SPIRIT-AI extension. Nature Medicine, 26(9), 1351–1363. https://doi.org/10.1038/s41591-020-1037-7

[37] Liu, X., Cruz Rivera, S., Moher, D., Calvert, M. J., Denniston, A. K., & SPIRIT-AI and CONSORT-AI Working Group. (2020). Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: The CONSORT-AI extension. Nature Medicine, 26(9), 1364–1374. https://doi.org/10.1038/s41591-020-1034-x

[38] Health Canada. (2025, January 30). Pan-Canadian AI for health (AI4H) guiding principles. Government of Canada. https://www.canada.ca/en/health-canada/corporate/transparency/health-agreements/pan-canadian-ai-guiding-principles.html

[39] Lewis, J. E. (Ed.). (2020). Indigenous protocol and artificial intelligence position paper. The Initiative for Indigenous Futures & the Canadian Institute for Advanced Research (CIFAR). https://doi.org/10.11573/spectrum.library.concordia.ca.00986506

[40] Carroll, S. R., Garba, I., Figueroa-Rodríguez, O. L., Holbrook, J., Lovett, R., Materechera, S., Parsons, M., Raseroka, K., Rodriguez-Lonebear, D., Rowe, R., Sara, R., Walker, J. D., Anderson, J., & Hudson, M. (2020). The CARE principles for Indigenous data governance. Data Science Journal, 19, Article 43. https://doi.org/10.5334/dsj-2020-043

[41] Government of Canada. (n.d.). Algorithmic Impact Assessment [online tool]. Retrieved September 28, 2026, from https://open.canada.ca/aia-eia-js/

[42] Tabassi, E. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.100-1

[43] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), 2024 O.J. (L 2024/1689). https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng

[44] International Medical Device Regulators Forum. (2014). Software as a medical device: Possible framework for risk categorization and corresponding considerations (IMDRF/SaMD WG/N12FINAL:2014). https://www.imdrf.org/documents/software-medical-device-possible-framework-risk-categorization-and-corresponding-considerations

[45] Strong, J., Rogers, H., Sun, E., Todsen, A. L., Ede, J., Lumley, C., Yeung, N., Higham, H., & Noble, J. A. (2026). Human-AI collaboration in healthcare: A scoping review. npj Digital Medicine. Advance online publication. https://doi.org/10.1038/s41746-026-02918-6

[46] Vaccaro, M., Almaatouq, A., & Malone, T. (2024). When combinations of humans and AI are useful: A systematic review and meta-analysis. Nature Human Behaviour, 8(12), 2293–2303. https://doi.org/10.1038/s41562-024-02024-1

[47] International Organization for Standardization. (2023). Information technology—Artificial intelligence—Management system (ISO/IEC Standard No. 42001:2023). https://www.iso.org/standard/42001

[48] Collins, G. S., Moons, K. G. M., Dhiman, P., Riley, R. D., Beam, A. L., Van Calster, B., Ghassemi, M., Liu, X., Reitsma, J. B., van Smeden, M., Boulesteix, A.-L., Camaradou, J. C., Celi, L. A., Denaxas, S., Denniston, A. K., Glocker, B., Golub, R. M., Harvey, H., Heinze, G., … Logullo, P. (2024). TRIPOD+AI statement: Updated guidance for reporting clinical prediction models that use regression or machine learning methods. BMJ, 385, e078378. https://doi.org/10.1136/bmj-2023-078378

[49] Verma, A. A., Stukel, T. A., Colacci, M., Bell, S., Ailon, J., Friedrich, J. O., Murray, J., Kuzulugil, S., Yang, Z., Lee, Y., Pou-Prom, C., & Mamdani, M. (2024). Clinical evaluation of a machine learning–based early warning system for patient deterioration. Canadian Medical Association Journal, 196(30), E1027–E1037. https://doi.org/10.1503/cmaj.240132

[50] Miao, B. Y., Chen, I. Y., Williams, C. Y. K., Davidson, J., Garcia-Agundez, A., Sun, S., Zack, T., Saria, S., Arnaout, R., Quer, G., Sadaei, H. J., Torkamani, A., Beaulieu-Jones, B., Yu, B., Gianfrancesco, M., Butte, A. J., Norgeot, B., & Sushil, M. (2025). The MI-CLAIM-GEN checklist for generative artificial intelligence in health. Nature Medicine, 31(5), 1394–1398. https://doi.org/10.1038/s41591-024-03470-0

[51] NHS England. (2026, July 29). Guidance on the use of AI-enabled ambient scribing products in health and care settings (Version 3). Retrieved September 28, 2026, from https://www.england.nhs.uk/long-read/guidance-on-the-use-of-ai-enabled-ambient-scribing-products-in-health-and-care-settings/

[52] Duggan, M. J., Gervase, J., Schoenbaum, A., Hanson, W., Howell, J. T., III, Sheinberg, M., & Johnson, K. B. (2025). Clinician experiences with ambient scribe technology to assist with documentation burden and efficiency. JAMA Network Open, 8(2), e2460637. https://doi.org/10.1001/jamanetworkopen.2024.60637

[53] Agarwal, P., Lall, R., & Girdhari, R. (2024). Artificial intelligence scribes in primary care. Canadian Medical Association Journal, 196(30), E1042. https://doi.org/10.1503/cmaj.240363

[54] Acharya, D. B., Kuppan, K., & Divya, B. (2025). Agentic AI: Autonomous intelligence for complex goals—A comprehensive survey. IEEE Access, 13, 18912–18936. https://doi.org/10.1109/ACCESS.2025.3532853

[55] Karunanayake, N. (2025). Next-generation agentic AI for transforming healthcare. Informatics and Health, 2(2), 73–83. https://doi.org/10.1016/j.infoh.2025.03.001

[56] Lin, G., Jiang, R., Elhadad, N., & Xu, X. (2026). A framework for longitudinal health AI agents. Nature Health, 1(8), 796–805. https://doi.org/10.1038/s44360-026-00141-x

[57] Kasirzadeh, A., & Gabriel, I. (2026). Agentic profiles for effective AI governance. Nature, 656(8127), 320–328. https://doi.org/10.1038/s41586-026-10805-z

[58] World Health Organization Regional Office for Europe. (2025). Health data governance in the age of artificial intelligence: Policy imperatives for the WHO European Region. https://www.who.int/europe/publications/i/item/WHO-EURO-2025-11462-51234-78079

[59] Organisation for Economic Co-operation and Development. (2022). Health data governance for the digital age: Implementing the OECD Recommendation on Health Data Governance. OECD Publishing. https://doi.org/10.1787/68b60796-en

[60] Pascoe, C., Quinn, S., & Scarfone, K. (2024). The NIST Cybersecurity Framework (CSF) 2.0 (NIST CSWP 29). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.CSWP.29

[61] Vassilev, A., Oprea, A., Fordyce, A., Anderson, H., Davies, X., & Hamin, M. (2025). Adversarial machine learning: A taxonomy and terminology of attacks and mitigations (NIST AI 100-2e2025). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.100-2e2025

[62] World Health Organization Regional Office for Europe. (2025). Cybersecurity and privacy maturity assessment and strengthening for digital health information systems. https://www.who.int/europe/publications/i/item/WHO-EURO-2025-11827-51599-78854

[63] European Commission. (2026, July 27). AI literacy: Questions & answers. Shaping Europe's Digital Future. https://digital-strategy.ec.europa.eu/en/faqs/ai-literacy-questions-answers

[64] Keren, R., Desai, B. R., & West, D. C. (2026). Promoting clinical expertise in the age of AI: No struggle, no mastery. JAMA, 335(21), 1837–1838. https://doi.org/10.1001/jama.2026.6097

[65] Ke, Y., Jin, L., Ong, J. C. L., Thirunavukarasu, A. J., Car, J., Cheung, C. Y., Tham, Y. C., Ting, D. S. W., Ong, M. E. H., Compton, S., Narayan, A., Keane, P. A., Wong, T. Y., Bates, D. W., Tan, P., & Liu, N. (2026). AI-induced never-skilling in medical education. Nature Medicine, 32(6), 1997–2006. https://doi.org/10.1038/s41591-026-04438-y

[66] Budzyń, K., Romańczyk, M., Kitala, D., Kołodziej, P., Bugajski, M., Adami, H. O., Blom, J., Buszkiewicz, M., Halvorsen, N., Hassan, C., Romańczyk, T., Holme, Ø., Jarus, K., Fielding, S., Kunar, M., Pellise, M., Pilonis, N., Kamiński, M. F., Kalager, M., … Mori, Y. (2025). Endoscopist deskilling risk after exposure to artificial intelligence in colonoscopy: A multicentre, observational study. The Lancet Gastroenterology & Hepatology, 10(10), 896–903. https://doi.org/10.1016/s2468-1253(25)00133-5

[67] Correction to Lancet Gastroenterol Hepatol 2025; 10: 896–903. (2025). The Lancet Gastroenterology & Hepatology, 10(11), e12. https://doi.org/10.1016/s2468-1253(25)00294-8

[68] Almyranti, M., Sutherland, E., Ash, N., & Eiszele, S. (2024). Artificial intelligence and the health workforce: Perspectives from medical associations on AI in health (OECD Artificial Intelligence Papers No. 28). OECD Publishing. https://doi.org/10.1787/9a31d8af-en

[69] Organisation for Economic Co-operation and Development. (2024). Using AI in the workplace: Opportunities, risks and policy responses (OECD Artificial Intelligence Papers No. 11). OECD Publishing. https://doi.org/10.1787/73d417f9-en

[70] Organisation for Economic Co-operation and Development. (2025). How widespread is algorithmic management in workplaces? OECD Publishing. https://doi.org/10.1787/cda7a114-en

[71] Temerty Centre for Artificial Intelligence Research and Education in Medicine. (n.d.). T-CAIREM AI Learning Hub. University of Toronto. Retrieved September 28, 2026, from https://learning.tcairem.ai/

[72] Temerty Centre for Artificial Intelligence Research and Education in Medicine. (n.d.). Start here. T-CAIREM AI Learning Hub. Retrieved September 28, 2026, from https://learning.tcairem.ai/start-here/

[73] Temerty Centre for Artificial Intelligence Research and Education in Medicine. (n.d.). Terminology. T-CAIREM AI Learning Hub. Retrieved September 28, 2026, from https://learning.tcairem.ai/terminology/

[74] Temerty Centre for Artificial Intelligence Research and Education in Medicine. (n.d.). International AI in Medicine Education Working Group. University of Toronto. Retrieved September 28, 2026, from https://tcairem.utoronto.ca/international-ai-medicine-education-working-group

[75] International AI in Medicine Education Working Group. (2025, September). Global framework for an AI curriculum in medicine. Temerty Centre for Artificial Intelligence Research and Education in Medicine. https://tcairem.utoronto.ca/sites/default/files/inline-files/working-group-framework-new-citations-19mar26%20conv.pdf

[76] T-CAIREM. (2026). T-CAIREM's AI educational offerings for medical residents. University of Toronto. Program information and enquiries: https://tcairem.utoronto.ca/t-cairem-post-graduate-medical-education

[77] Rosella, L. C., Aphinyanaphongs, Y., Barry, J., Boyko, O., Castro, H., Chang, A., Eaton, S., Erickson, B. J., Fairclough, J., Ferro, D., Gascon, G., Haque, N., Hoyt, R., Kim, J., Limon, A., Lio, P., Mamdani, M., Mathur, P., McAdams, R., … Yung, N. (2026). A global framework for artificial intelligence education in medicine: International

working group recommendations. npj Digital Medicine. Advance online publication. https://doi.org/10.1038/s41746-026-03197-x

[78] Pasquini, G., Stocking, G., Kikuchi, E., Pula, I., & Yam, E. (2026, April 7). Users of social media and AI chatbots for health information are more likely to say they are convenient than accurate. Pew Research Center. https://www.pewresearch.org/science/2026/04/07/users-of-social-media-and-ai-chatbots-for-health-information-are-more-likely-to-say-they-are-convenient-than-accurate/

[79] World Health Organization. (2022, February 9). Ensuring artificial intelligence (AI) technologies for health benefit older people. https://www.who.int/news/item/09-02-2022-ensuring-artificial-intelligence-(ai)-technologies-for-health-benefit-older-people

[80] Bean, A. M., Payne, R. E., Parsons, G., Kirk, H. R., Ciro, J., Mosquera-Gómez, R., Hincapié M, S., Ekanayaka, A. S., Tarassenko, L., Rocher, L., & Mahdi, A. (2026). Reliability of LLMs as medical assistants for the general public: A randomized preregistered study. Nature Medicine, 32(2), 609–615. https://doi.org/10.1038/s41591-025-04074-y. Publisher correction published April 17, 2026: https://doi.org/10.1038/s41591-026-04404-8

[81] Sujan, M., Pool, R., & Salmon, P. (2022). Eight human factors and ergonomics principles for healthcare artificial intelligence. BMJ Health & Care Informatics, 29(1), e100516. https://doi.org/10.1136/bmjhci-2021-100516

[82] Coalition for Health AI. (2026, May 27). AI governance playbooks. Retrieved September 28, 2026, from https://www.chai.org/workgroup/cross-cutting/ai-governance

[83] Lawrence, K., Kuram, V. S., Levine, D. L., Sharif, S., Polet, C., Malhotra, K., & Owens, K. (2025). Informed consent for ambient documentation using generative AI in ambulatory care. JAMA Network Open, 8(7), e2522400. https://doi.org/10.1001/jamanetworkopen.2025.22400

[84] Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. Science, 366(6464), 447–453. https://doi.org/10.1126/science.aax2342

[85] Wang, H. E., Landers, M., Adams, R., Subbaswamy, A., Kharrazi, H., Gaskin, D. J., & Saria, S. (2022). A bias evaluation checklist for predictive models and its pilot application for 30-day hospital readmission models. Journal of the American Medical Informatics Association, 29(8), 1323–1333. https://doi.org/10.1093/jamia/ocac065

[86] Global Indigenous Data Alliance. (n.d.). CARE principles for Indigenous data governance. Retrieved September 28, 2026, from https://www.gida-global.org/careprinciples

[87] First Nations Information Governance Centre. (n.d.). The First Nations principles of OCAP®. Retrieved September 28, 2026, from https://fnigc.ca/ocap-training/ OCAP® is a registered trademark of the First Nations Information Governance Centre (FNIGC).

[88] Whittaker, R., Dobson, R., Jin, C. K., Style, R., Jayathissa, P., Hiini, K., Ross, K., Kawamura, K., Muir, P., & Waitematā AI Governance Group. (2023). An example of governance for AI in health services from Aotearoa New Zealand. npj Digital Medicine, 6(1), 164. https://doi.org/10.1038/s41746-023-00882-z

[89] Matos, J., Van Calster, B., Celi, L. A., Dhiman, P., Gichoya, J. W., Riley, R. D., Russell, C., Khalid, S., & Collins, G. S. (2026). Critical appraisal of fairness metrics for artificial intelligence-based clinical prediction models: A scoping review. The Lancet Digital Health, 8(7), 101001. https://doi.org/10.1016/j.landig.2026.101001

[90] Smith, C., Boyd, A., Pagel, C., & Morris, A. (2026). Trust, not technology: Governing access to health data as the decisive challenge for the UK. The Lancet Digital Health, 8(5), 100994. https://doi.org/10.1016/j.landig.2026.100994

[91] Rao, A. K., Keller, A. J., Kalra, N., Steed, R., Kwegyir-Aggrey, K., Klyman, K., Staheli, D., & Bergman, A. S. (2026). Challenges to the monitoring of deployed AI systems: Center for AI Standards and Innovation (NIST AI 800-4). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.800-4

[92] International Organization for Standardization. (2025). Information technology—Artificial intelligence (AI)—AI system impact assessment (ISO/IEC Standard No. 42005:2025). https://www.iso.org/standard/42005

[93] Lekadir, K., Frangi, A. F., Porras, A. R., Glocker, B., Cintas, C., Langlotz, C. P., Weicken, E., Asselbergs, F. W., Prior, F., Collins, G. S., Kaissis, G., Tsakou, G., Buvat, I., Kalpathy-Cramer, J., Mongan, J., Schnabel, J. A., Kushibar, K.,

Riklund, K., Marias, K., … Starmans, M. P. A. (2025). FUTURE-AI: International consensus guideline for trustworthy and deployable artificial intelligence in healthcare. BMJ, 388, e081554. https://doi.org/10.1136/bmj-2024-081554

[94] Nong, P., Hamasha, R., Singh, K., Adler-Milstein, J., & Platt, J. (2024). How academic medical centers govern AI prediction tools in the context of uncertainty and evolving regulation. NEJM AI, 1(3). https://doi.org/10.1056/aip2300048

[95] Liao, F., Adelaine, S., Afshar, M., & Patterson, B. W. (2022). Governance of clinical AI applications to facilitate safe and equitable deployment in a large health system: Key elements and early successes. Frontiers in Digital Health, 4, 931439. https://doi.org/10.3389/fdgth.2022.931439

[96] Callahan, A., McElfresh, D., Banda, J. M., Bunney, G., Char, D., Chen, J., Corbin, C. K., Dash, D., Downing, N. L., Jain, S. S., Kotecha, N., Masterson, J., Mello, M. M., Morse, K., Nallan, S., Pandya, A., Revri, A., Sharma, A., Sharp, C., … Shah, N. H. (2024). Standing on FURM ground: A framework for evaluating fair, useful, and reliable AI models in health care systems. NEJM Catalyst, 5(10). https://doi.org/10.1056/cat.24.0131

[97] Kim, J. Y., Hasan, A., Kueper, J., Tang, T., Hayes, C., Fine, B., Balu, S., & Sendak, M. (2025). Establishing organizational AI governance in healthcare: A case study in Canada. npj Digital Medicine, 8(1), 522. https://doi.org/10.1038/s41746-025-01909-3

[98] Kim, J. Y., Hasan, A., Balu, S., & Sendak, M. (2026). People process technology and operations framework for establishing AI governance in healthcare organizations. npj Digital Medicine, 9(1), 224. https://doi.org/10.1038/s41746-026-02419-6

[99] NHS Digital. (2018). DCB0160: Clinical risk management—Its application in the deployment and use of health IT systems (Amd 25/2018). https://digital.nhs.uk/data-and-information/information-standards/governance/latest-activity/standards-and-collections/dcb0160-clinical-risk-management-its-application-in-the-deployment-and-use-of-health-it-systems

[100] Zai, A. H., Adibuzzaman, M., McManus, D. D., & Walkey, A. (2026). The HALO model: A learning health system framework for artificial intelligence. Learning Health Systems, 10(3), e70100. https://doi.org/10.1002/lrh2.70100

[101] National Institute of Standards and Technology. (n.d.). AI Risk Management Framework. Retrieved October 1, 2026, from https://www.nist.gov/itl/ai-risk-management-framework

[102] National Institute of Standards and Technology. (2026, June 10). NIST AI RMF playbook. Retrieved October 1, 2026, from https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook

[103] Health Canada. (2026, April 1). Pre-market guidance for machine learning-enabled medical devices. Government of Canada. https://www.canada.ca/en/health-canada/services/drugs-health-products/medical-devices/application-information/guidance-documents/pre-market-guidance-machine-learning-enabled-medical-devices.html

[104] Joint Commission. (2026, June 1). Joint Commission releases first of its kind exclusively designed for healthcare organizations, voluntary responsible use of AI in healthcare certification. https://www.jointcommission.org/en-us/knowledge-library/news/2026-05-responsible-use-of-ai-in-healthcare-certification

[105] NHS England. (2026, June 30). Review of digital clinical safety standards: DCB0129 and DCB0160. NHS England Digital. Retrieved October 1, 2026, from https://digital.nhs.uk/data-and-information/information-standards/governance/latest-activity/standards-and-collections/review-of-digital-clinical-safety-standards-dcb0129-and-dcb0160

[106] European Commission. (2026, August 7). Navigating the AI Act. Shaping Europe's Digital Future. Retrieved October 1, 2026, from https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act

[107] U.S. Food and Drug Administration. (2026, August 3). Applying human factors and usability engineering to medical devices: Guidance for industry and Food and Drug Administration staff. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/applying-human-factors-and-usability-engineering-medical-devices

[108] U.S. Food and Drug Administration. (2026, January 29). Clinical decision support software: Guidance for industry and Food and Drug Administration staff. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software

[109] Sounderajah, V., Guni, A., Liu, X., Collins, G. S., Karthikesalingam, A., Markar, S. R., Golub, R. M., Denniston, A. K., Shetty, S., Moher, D., Bossuyt, P. M., Darzi, A., Ashrafian, H., & STARD-AI Steering Committee. (2025). The STARD-AI reporting guideline for diagnostic accuracy studies using artificial intelligence. Nature Medicine, 31(10), 3283–3289. https://doi.org/10.1038/s41591-025-03953-8

[110] Moons, K. G. M., Damen, J. A. A., Kaul, T., Hooft, L., Andaur Navarro, C., Dhiman, P., Beam, A. L., Van Calster, B., Celi, L. A., Denaxas, S., Denniston, A. K., Ghassemi, M., Heinze, G., Kengne, A. P., Maier-Hein, L., Liu, X., Logullo, P., McCradden, M. D., Liu, N., … van Smeden, M. (2025). PROBAST+AI: An updated quality, risk of bias, and applicability assessment tool for prediction models using regression or artificial intelligence methods. BMJ, 388, e082505. https://doi.org/10.1136/bmj-2024-082505

[111] Wekenborg, M. K., Gilbert, S., & Kather, J. N. (2025). Examining human-AI interaction in real-world healthcare beyond the laboratory. npj Digital Medicine, 8(1), 169. https://doi.org/10.1038/s41746-025-01559-5

[112] Biro, J. M., Handley, J. L., McCurry, J. M., Visconti, A., Weinfeld, J., Trafton, J. G., & Ratwani, R. M. (2025). Opportunities and risks of artificial intelligence in patient portal messaging in primary care. npj Digital Medicine, 8(1), 222. https://doi.org/10.1038/s41746-025-01586-2

[113] Ma, S. P., Liang, A. S., Shah, S. J., Smith, M., Jeong, Y., Devon-Sand, A., Crowell, T., Delahaie, C., Hsia, C., Lin, S., Shanafelt, T., Pfeffer, M. A., Sharp, C., & Garcia, P. (2025). Ambient artificial intelligence scribes: Utilization and impact on documentation time. Journal of the American Medical Informatics Association, 32(2), 381–385. https://doi.org/10.1093/jamia/ocae304

[114] Olson, K. D., Meeker, D., Troup, M., Barker, T. D., Nguyen, V. H., Manders, J. B., Stults, C. D., Jones, V. G., Shah, S. D., Shah, T., & Schwamm, L. H. (2025). Use of ambient AI scribes to reduce administrative burden and professional burnout. JAMA Network Open, 8(10), e2534976. https://doi.org/10.1001/jamanetworkopen.2025.34976

[115] Kanaparthy, N. S., Villuendas-Rey, Y., Bakare, T., Diao, Z., Iscoe, M., Loza, A., Wright, D., Safranek, C., Faustino, I. V., Brackett, A., Melnick, E. R., & Taylor, R. A. (2025). Real-world evidence synthesis of digital scribes using ambient listening and generative artificial intelligence for clinician documentation workflows: Rapid review. JMIR AI, 4, e76743. https://doi.org/10.2196/76743

[116] Woznitza, N., Smith, L., Rawlinson, J., Au-Yong, I., George, B., Djearaman, M. G., Nair, A., Lee, R. W., Navani, N., Ndwandwe, S., Clarke, C. S., Creeden, A., Newsome, J., Das, I., Abaokporo, S., Tucker, R., Hathorn, J., & Baldwin, D. R. (2026). AI-based chest X-ray prioritization in the lung cancer diagnostic pathway: The LungIMPACT randomized controlled trial. Nature Medicine, 32(5), 1737–1744. https://doi.org/10.1038/s41591-026-04253-5

[117] Agweyu, A., Mwaniki, P., Menon, V., Korom, R., Isaaka, L., Wanyama, C., Gill, J., Kiptinness, S., Adan, N., Emmanuel-Fabula, M., Riley, R. D., Archer, L., Denniston, A. K., Liu, X., & Mateen, B. A. (2026). Generative AI-enabled clinical decision support system in primary care: A pragmatic, cluster-randomized trial. Nature Medicine, 32(8), 3032–3039. https://doi.org/10.1038/s41591-026-04503-6

[118] Lukac, P. J., Turner, W., Vangala, S., Chin, A. T., Khalili, J., Shih, Y.-C. T., Sarkisian, C., Cheng, E. M., & Mafi, J. N. (2025). Ambient AI scribes in clinical practice: A randomized trial. NEJM AI, 2(12), AIoa2501000. https://doi.org/10.1056/AIoa2501000

[119] Office of the Privacy Commissioner of Canada. (2023, December 7). Principles for responsible, trustworthy and privacy-protective generative AI technologies. https://www.priv.gc.ca/en/privacy-topics/technology/artificial-intelligence/gd_principles_ai/

[120] Temerty Centre for Artificial Intelligence Research and Education in Medicine. (n.d.). Education overview. University of Toronto. Retrieved October 1, 2026, from https://tcairem.utoronto.ca/education-overview-0

[121] Temerty Centre for Artificial Intelligence Research and Education in Medicine. (n.d.). T-CAIREM & post-graduate medical education. University of Toronto. Retrieved October 1, 2026, from https://tcairem.utoronto.ca/t-cairem-post-graduate-medical-education

[122] Temerty Centre for Artificial Intelligence Research and Education in Medicine. (n.d.). AI in medicine upskilling courses. University of Toronto. Retrieved October 1, 2026, from https://tcairem.utoronto.ca/ai-medicine-upskilling-courses

[123] Temerty Centre for Artificial Intelligence Research and Education in Medicine. (n.d.). Education partnerships. University of Toronto. Retrieved October 1, 2026, from https://tcairem.utoronto.ca/education-partnerships

[124] T-CAIREM. (n.d.). EducationResources [GitHub repository]. GitHub. Retrieved October 1, 2026, from https://github.com/T-CAIREM/EducationResources

[125] Temerty Centre for Artificial Intelligence Research and Education in Medicine. (n.d.). Temerty Centre Speaker Series. University of Toronto. Retrieved October 1, 2026, from https://tcairem.utoronto.ca/temerty-centre-speaker-series

[126] Temerty Centre for Artificial Intelligence Research and Education in Medicine. (n.d.). T-CAIREM–TAHSN AI in healthcare (TAIH) Community of Practice. University of Toronto. Retrieved October 1, 2026, from https://tcairem.utoronto.ca/t-cairem-tahsn-al-healthcare-taih-community-practice

[127] Temerty Centre for Artificial Intelligence Research and Education in Medicine. (2025, November 18). AI research in healthcare: Workshop for patient partners (virtual). University of Toronto. https://tcairem.utoronto.ca/event/ai-research-healthcare-workshop-patient-partners-virtual

[128] Innovation, Science and Economic Development Canada. (2026, June 1). AI Sovereign Compute Infrastructure Program. https://ised-isde.canada.ca/site/ised/en/ai-sovereign-compute-infrastructure-program

[129] Feng, J., Phillips, R. V., Malenica, I., Bishara, A., Hubbard, A. E., Celi, L. A., & Pirracchio, R. (2022). Clinical artificial intelligence quality improvement: Towards continual monitoring and updating of AI algorithms in healthcare. npj Digital Medicine, 5, 66. https://doi.org/10.1038/s41746-022-00611-y

[130] Castro, D. C., Walker, I., & Glocker, B. (2020). Causality matters in medical imaging. Nature Communications, 11, 3673. https://doi.org/10.1038/s41467-020-17478-w

[131] Kore, A., Abbasi Bavil, E., Subasri, V., Abdalla, M., Fine, B., Dolatabadi, E., & Abdalla, M. (2024). Empirical data drift detection experiments on real-world medical imaging data. Nature Communications, 15, 1887. https://doi.org/10.1038/s41467-024-46142-w

[132] Ong Ly, C., Unnikrishnan, B., Tadic, T., Patel, T., Duhamel, J., Kandel, S., Moayedi, Y., Brudno, M., Hope, A., Ross, H., & McIntosh, C. (2024). Shortcut learning in medical AI hinders generalization: Method for estimating AI model generalization without external data. npj Digital Medicine, 7, 124. https://doi.org/10.1038/s41746-024-01118-4

[133] Bernhardt, M., Castro, D. C., Tanno, R., Schwaighofer, A., Tezcan, K. C., Monteiro, M., Bannur, S., Lungren, M. P., Nori, A., Glocker, B., Alvarez-Valle, J., & Oktay, O. (2022). Active label cleaning for improved dataset quality under resource constraints. Nature Communications, 13, 1161. https://doi.org/10.1038/s41467-022-28818-3

[134] Finlayson, S. G., Subbaswamy, A., Singh, K., Bowers, J., Kupke, A., Zittrain, J., Kohane, I. S., & Saria, S. (2021). The clinician and dataset shift in artificial intelligence. New England Journal of Medicine, 385(3), 283–286. https://doi.org/10.1056/NEJMc2104626

[135] Temoshok, D., Fenton, J. L., Choong, Y.-Y., Lefkovitz, N., Regenscheid, A., Galluzzo, R., & Richer, J. P. (2025). Digital identity guidelines: Authentication and authenticator management (NIST SP 800-63B-4). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-63B-4

[136] Geirhos, R., Jacobsen, J.-H., Michaelis, C., Zemel, R., Brendel, W., Bethge, M., & Wichmann, F. A. (2020). Shortcut learning in deep neural networks. Nature Machine Intelligence, 2, 665–673. https://doi.org/10.1038/s42256-020-00257-z

[137] Krakovna, V., Uesato, J., Mikulik, V., Rahtz, M., Everitt, T., Kumar, R., Kenton, Z., Leike, J., & Legg, S. (2020, April 21). Specification gaming: The flip side of AI ingenuity. DeepMind. https://deepmind.google/blog/specification-gaming-the-flip-side-of-ai-ingenuity/

[138] Von Arx, S., Chan, L., & Barnes, B. (2025, June 5). Recent frontier models are reward hacking. METR. https://metr.org/blog/2025-06-05-recent-reward-hacking/

[139] OpenAI. (2026, August 26). OpenAI–Hugging Face incident technical report. https://cdn.openai.com/pdf/67869394-cb91-4c12-888c-5cbd85c7814c/OpenAI-Hugging-Face%20Incident-Technical-Report.pdf

[140] Greenblatt, R., Cotra, A., & Wijk, H. (2026, August 26). Brief independent investigation of agents' behavior, reasoning and collaboration in the OpenAI / Hugging Face hacking incident. METR. https://metr.org/blog/2026-08-26-openai-hugging-face-incident-investigation/

Appendix A Worked examples

The following cases show how a hospital could review a request to expand an AI service that is already approved and in use. Each follows the decision from the initial request through evaluation and release, then considers monitoring, incidents, and retirement. The hospital may decide to continue, expand, restrict, or retire the service according to the evidence and circumstances described. All decisions are hypothetical; the cases report no actual approval, measured performance, or implementation outcome. A hospital using them must make its own findings under applicable law and professional requirements and consider whether it has the staff, resources, and authority to accept the remaining risk. The examples prescribe no numerical thresholds or sample sizes and assume no satisfactory evaluation results.

Each case uses the service's existing registry entry and dossier. The registry shows what is operating and what the hospital has authorized; the linked dossier holds the supporting evidence and decision history. Reviewers compare the expansion request with that approval, keeping one consistent account of the service. The clinical/end-user owner, internal technical owner, and operational owner retain the responsibilities assigned in **Section 7** throughout review and operation. Their work differs by service, but approval does not end their responsibilities.

All three cases provisionally fall in Tier 3 because errors could have serious consequences, but an authorized person is expected to review the output meaningfully before a consequential action. Reviewers must verify both that this limit works and that staff have enough time and resources to carry out the review. If either assumption fails, they reconsider Tier 4 and reopen the relevant Tier 0 questions about whether the activity may proceed. All three cases require full review. **Table 1** identifies the seven decision stages and the approvals needed at each. The six recurring governance tasks support those decisions; the numbered steps in each case organize the example and do not add lifecycle stages.

The cases put the AI literacy expectations in **Table 3** into practice. Owners and users demonstrate the knowledge and skills their roles require, and the hospital considers what patients need to understand when the service affects them. These assessments belong in the same review, using the common record fields in **Appendices I and L** rather than a separate case file.

A1 Inpatient deterioration alert

In this example, two adult general-medicine units use a specified model release to identify patients at increased risk of deterioration. The model calculates risk scores from approved electronic health record (EHR) data at a defined interval. An approved threshold determines which scores trigger an alert; this is the classification function in this case. An alert goes to the assigned nurse and medical team, who begin a bedside assessment under the hospital's existing escalation protocol. The alert does not diagnose the patient or establish that the patient is stable. It cannot independently place orders, change monitoring, or transfer care. Staff must continue observing patients and escalating concerns whether or not an alert appears.

The hospital proposes extending the alert to a third adult general-medicine unit. Reviewers must decide whether the existing evidence applies to its patients and workflow and whether its staff can manage the expected alerts and clinical responses. They must also decide whether the third unit can remain in the same configuration group: a set of variants that the hospital can evaluate, monitor, and support together within approved limits. Here, the configuration includes the model version, settings, data connections, alert delivery, and clinical response. Reviewers must show that differences between units do not undermine the shared evidence or monitoring and support arrangements; using the same product is not enough. They record the comparison and full expansion decision in the existing registry entry and dossier, following **Table 1**'s seven decision stages and **Figure 5**'s six recurring tasks. [15][34]

A1.1 Define the change and decision authority

The governance lead reopens the record for an expansion decision and compares the current service with the proposed third-unit setup. For each, the record identifies the model version, alert threshold, eligibility and scoring rules, source-data route, alert display and routing, and clinical response protocol. It names the current and proposed units, expected volume, and maximum scale the hospital would authorize. The lead obtains current performance and capacity evidence from the first two units, including incidents and unresolved conditions, before the expansion decision. The third unit cannot begin live use until expansion is authorized and a separate pre-release check—production confirmation—verifies that its actual setup meets the approval conditions. As **Figure 1** shows, approval covers this deterioration-alert use, not every function of the product or model.

The three owners confirm their **Section 7** responsibilities for the additional unit. The clinical owner is accountable for clinical criteria, the response protocol, and consequences of use. The internal technical owner is accountable for the version in use, traceable data history, delivery to the right staff, monitoring, and restoration of an earlier approved setup if needed (rollback). The operational owner establishes whether staffing, competence, support, and the ordinary clinical pathway can cover the additional unit. The record names who performs each check, accepts its evidence, and responds to a problem. The AI Steering Committee uses these findings to decide within its delegated authority and risk limits, referring decisions beyond those limits to the authorized higher body. Specialists retain authority over findings within their mandates; the committee cannot waive a mandatory finding. Suppliers and partners may provide evidence or correct faults, but the hospital decides whether to expand, restart, or retire the service.

The governance lead applies Tier 0 to the expansion review and any later evaluation. Reject means the proposal is prohibited or cannot be remedied; Hold means a prerequisite is unresolved; Proceed permits only the specified next activity. Reviewers need a clear purpose, named owners, and a data route they can assess. The record states the decision, reason, conditions, any corrective action, and who may reconsider it. Privacy, cybersecurity, clinical safety, and other required specialist findings must be obtained before the data access, system connection, or evaluation activity to which they apply.

Some checks need to be completed before live use rather than before preliminary review. Before release, the hospital must verify that records can show how an alert was produced and delivered, that ordinary clinical escalation remains available, and that it can suppress alerts or restore an earlier approved setup. Earlier review may continue while these checks are underway only if that work does not depend on an unresolved control. A Proceed decision does not authorize live use.

Every unresolved specialist or supplier requirement must be tied to the activity it prevents. The governance lead records its owner, due date, evidence needed to resolve it, and the person authorized to accept that evidence. For example, missing information about the release may prevent local evaluation, while an unverified link between live source fields and model inputs may block release even after evaluation is complete. The affected activity remains on hold until the requirement is met, unless an authorized decision narrows or redesigns the proposal or revises the contract to resolve it. Unrelated preparatory work may continue only within its recorded permission. Assigning a lower consequence tier cannot compensate for missing evidence or controls.

Reviewers assess the most serious harm that could credibly occur within the service's verified operating limits. They do not assume planned safeguards already work. A missing, delayed, or misleading alert could delay recognition of deterioration if staff rely on it too heavily. False alerts could divert attention, cause unnecessary escalation, or produce alert fatigue. Incorrect routing could disclose information to the wrong team or leave the responsible team unaware that assessment is needed. For each route to harm, the record identifies the people affected, how widely they could be exposed, how quickly harm could occur, whether staff are likely to detect it in time, and whether its effects can be reversed. Reviewers also identify groups for whom eligibility, data completeness, or performance may differ.

Reviewers first ask whether Tier 4 applies. Tier 3 is justified only if an authorized clinician can assess the patient before a consequential action and has the information, time, competence, and authority to do so. The clinician must be able to disregard an inappropriate alert and continue independent observation and escalation when alerts are missing or delayed. Tier 4 must be reconsidered if the service changes monitoring, places orders, initiates transfers, determines clinically significant access or priority, or suppresses review. It must also be reconsidered if serious harm could occur before a meaningful clinical decision, or if workload or interface design makes approval a formality. Effective safeguards may reduce the risk that remains, but the credible Tier 3 consequences still determine the classification. This service requires full review and cannot use the streamlined lower-risk route.

As part of the expansion review, the owners check whether the existing service still matches its last approval. If the historical dossier is incomplete, they reconstruct the current setup from configuration records, operating logs, contracts, procedures, and observation of staff practice. An incomplete record does not automatically require withdrawal, which may itself create a safety risk. The owners first contain immediate hazards, then obtain a decision supported by evidence to continue within stated limits, restrict use, correct deficiencies, or retire the service. Expansion waits until the current setup and its authorization are established. If a stop condition is met, however, the service must be suspended and the ordinary clinical pathway used; missing historical records cannot justify continuing beyond that condition.

A1.2 Confirm need and use boundaries

The clinical and operational owners describe how staff on the proposed unit currently recognize deterioration and escalate concerns. Over a defined baseline period, they examine patient volume, assessment timing, rapid-response activity, and unplanned transfers. They also examine workload and known failures in that pathway. They compare the proposed alert with feasible alternatives, such as stronger observation practices, changes to rounding or escalation, staffing support, or an existing non-AI score. Availability of the model elsewhere in the hospital, or frequent acknowledgement of its alerts, does not establish that the third unit needs it.

Before collecting further evidence, the owners agree what improvement would justify expansion and how to measure it. The aim may be earlier bedside assessment, more useful warning time, or completion of indicated escalation. It may instead be to reduce alert burden or staff workload, or to improve selected patient outcomes. The evaluation also measures unnecessary tests, repeated alerts, interruptions, and work passed to other teams, so any benefit can be weighed against costs to care and staff time. Each measure specifies whose results are counted, the denominator used to calculate a rate, the observation period, the data source, and the result required for acceptance. Adoption and acknowledgement rates show whether the service is used; they cannot alone establish clinical benefit. **Appendices E and I** retain the measures and resulting decisions.

The approved-use statement explains which adult general-medicine patients are included or excluded, what counts as deterioration, and how far ahead the model predicts it. It specifies data sources, scoring interval, and alert threshold, then describes who receives the alert, how it appears, how staff may rely on it, and what they must do. It also states the alert's status in the clinical record and which ordinary clinical process applies if the service is unavailable or unsuitable. The response protocol names primary and backup recipients, assigns responsibility during handoffs, and states what staff must do if an alert is not acknowledged or the patient is not assessed within the clinically approved time.

Acknowledgement alone does not establish that assessment or indicated escalation occurred. Staff are told explicitly that no alert does not mean no deterioration and does not replace observation or escalation. Pediatrics, obstetrics, emergency care, and critical care are outside scope, as are use as a discharge criterion and autonomous orders or transfers. Introducing any of these uses requires a separate review or material-change review before use.

Patients and families help assess whether expansion would change communication, impose unequal burdens from escalation, or affect access to clinical attention. An alert shown to clinicians does not require encounter-specific disclosure in every setting. The responsible authority records whether disclosure applies and why. Where patients should receive information, the hospital identifies who will explain the service, what they will explain, and how patients can ask questions, request correction, or make a complaint.

A1.3 Assign hospital and supplier responsibilities

In this example, the hospital maintains the model, input-variable definitions, alert threshold, software release package, and evaluation code. Informatics maintains extraction from the EHR and delivery of alerts; the clinical service remains responsible for escalation. The EHR platform and some supporting infrastructure are purchased. The responsibility table described in **Figure 3** and **Appendix E** assigns work to hospital teams and suppliers, including failures that cross the EHR, scoring service, and clinical workflow. For each part, it identifies the required evidence, maintenance and response duties, and hospital owner who accepts the work.

For the components it controls, the hospital retains the data and model-development history, along with testing and release evidence. It maintains access controls, monitoring, and a process for correcting faults. Those duties continue through rollback, continuity, and retirement. When the hospital changes input variables, units, thresholds, suppression rules, mappings, or interfaces, it remains responsible for those changes. Buying the EHR platform does not transfer that responsibility to the supplier. The sponsor identifies the recurring funding and staff capacity needed to perform this work throughout operation. A funding commitment must cover the approved scale, including investigation and support when the service is restricted or paused.

Purchased components remain subject to **Appendix G**'s supplier requirements. These cover release and change evidence, cooperation during incidents, correction timelines, rollback or safe disabling, continuity, and exit. If an academic or regional partner validates or maintains a component, the responsibility table records its duties and the handoff of work and evidence to the hospital. It also identifies the contact, response deadline, and escalation when a party does not deliver. The hospital continues local monitoring and retains decisions to expand, restrict, pause, restart, or retire the service.

This example assumes that the hospital has approved enough staff and resources to support one adult general-medicine alert configuration group. The group uses a documented common setup: model version, input variables, threshold, scoring and suppression rules, display, routing, and response protocol. Unit names, rosters, and support contacts may differ if evidence shows those differences remain within the approved limits. A different threshold, recalibration, population, input mapping, interface, or response protocol needs a separate group unless reviewers establish that it can safely share the existing evidence, monitoring, and support. If the hospital already supports its maximum number of groups, it must consolidate or retire a group, or fund more capacity and obtain a new governance decision. Departments cannot introduce unregistered versions. The one-group limit is an assumption for this example and reflects the framework's default planning safeguard; it is neither an empirically established maximum nor a prescribed limit for every alert service.

A1.4 Review evidence for the new unit

The evidence reviewer brings together the clinical, technical, operational, and specialist findings. The review must establish whether the model predicts the intended event, whether the live data and scoring process preserve that performance, and whether staff can respond safely. Good model results alone cannot establish reliable data delivery or clinical response. The reviewer specifies which findings the committee needs to decide on expansion and which further checks must be completed before the new unit goes live. **Appendix D** sets out Tier 3 evidence requirements, **Appendix E** provides the lifecycle record, and **Appendix F** provides checks for prediction, classification, detection, and prioritization.

The first two units are assumed to have been selected for measured need, representative work, enforceable response standards, and sufficient support. Their actual operating evidence must be examined before expansion to the third unit is authorized. The evidence identifies the observation period, eligible population, and denominators for each rate. It also records missing or delayed evidence, the status of stop conditions, serious incidents, and corrective actions. This case does not assume satisfactory findings. The same definitions and records provide the starting point for monitoring the third unit after release.

Reviewers compare the third unit with the approved service before deciding which existing evidence can be reused. They examine patients and deterioration prevalence, data capture, workflow, staffing, shifts, escalation, accessibility needs, expected alert volume, and the staff and resources available for support and monitoring. These conditions may differ even if the patients appear similar. Evidence can be reused when the comparison shows that it still applies and that the new unit can use the same controls. Local checks of data flow and workflow remain necessary. Each later unit needs its own comparison, expansion decision, and production confirmation before live use.

The team compares each external study with the proposed model release, deterioration definition, prediction period, and threshold. It examines differences in the adult general-medicine population, prevalence, setting, and response workflow. A different patient mix or response to alerts may make a study informative without validating this service. The evidence inventory records those limits, missing subgroup evidence, conflicts, and developer involvement. It also records the reviewer's reason for accepting or rejecting each source.

Before the new results are known, the evidence reviewer records a local validation plan agreed with the clinical and technical owners. It states what the alert will be compared with and how reviewers will establish whether deterioration occurred (the reference standard). It defines the event and prediction periods, eligibility, sampling, and acceptance criteria. Evaluation uses the data extraction and scoring process intended for live operation and covers the proposed unit's days, shifts, and patient complexity. Cases come from the eligible population and the source used to identify outcomes, including patients who received no alert. The clinical reviewer specifies how outcomes will be assessed and disagreements resolved. Where feasible, the team evaluates new cases as they occur, generating scores without showing them to staff or using them in care. The data access and environment require the applicable authorization. If another method is used, the dossier explains it and what it can establish.

The report must make clear what each count represents: patients, encounters, deterioration events, or alerts. It explains how repeated alerts and overlapping prediction periods are handled, identifies eligible patients who were not scored, and gives the reasons. For every alert measure, the plan states whether it counts alerts generated, delivered, or acknowledged. Missing outcomes and incomplete follow-up are reported separately; they cannot be treated as evidence that deterioration did not occur. These counting rules prevent repeated alerts or missing cases from making performance appear better than it is.

Across the scored population, reviewers compare estimated risks with observed outcomes in clinically relevant risk ranges to assess calibration. At the approved alert threshold, they assess sensitivity: the proportion of deterioration events identified. They also assess positive predictive value: the proportion of alerts followed by the defined event. Other measures cover the burden of alerts not followed by that event (false positives), missed events (false negatives), and clinically useful lead time.

To measure sensitivity and find missed events, reviewers must identify deterioration independently of the alert list. For the whole service, the denominator is every deterioration event confirmed by the reference standard in the eligible population, including events in patients with no score or alert. Before evaluation, reviewers define how to match an event to an alert and how early that alert must arrive to support a clinically useful response. An alert generated but never delivered to the responsible team cannot count as successful detection by the live service. The report therefore separates model performance in cases with valid scores from detection achieved by the service as a whole. Positive predictive value uses the specified alerts or alert episodes as its denominator and

checks for the associated event within the agreed prediction period. Reviewing alerts alone cannot establish sensitivity or the rate of missed events.

The report separates deterioration after a below-threshold score from deterioration associated with no score, suppressed alerts, late delivery, or failed routing. Otherwise, an average model result could hide a failure in the service. Reviewers record care provided after alerts and explain how that care limits interpretation of later outcomes. They also distinguish discrimination—the model's ability to separate higher-risk from lower-risk patients—from clinical usefulness. They examine eligibility, missing or late data, and excluded or unscored patients separately. An acceptable overall result cannot replace the minimum subgroup results set before evaluation. The report makes small event counts and other uncertainty explicit and identifies the additional evidence or restrictions needed. A non-significant difference does not establish that groups have equal results. [48][18]

The technical team follows the complete alert process: source fields, score calculation, threshold and suppression rules, display, routing, acknowledgement, and clinical documentation. Tests include stale or missing values, changed codes or units, duplicate or late messages, and delivery to the wrong patient or team. They also include eligible patients with no score or displayed alert. The record timestamps the original observation, its receipt by the scoring service, score calculation, display, delivery, acknowledgement, and clinical assessment separately. Reviewers can then see how much warning time remains after data, delivery, and clinical-response delays. The team also tests downtime, recovery, and reconciliation. Correct calculations on historical data do not establish that the live service will deliver the right information to the right person in time.

The operational and clinical owners arrange workflow tests during conditions such as nights, weekends, busy handoffs, and interruptions. Representative staff work through false, missing, delayed, conflicting, and high-volume alerts. Reviewers observe bedside assessment and established escalation without treating a score as a diagnosis. Tests include deterioration without an alert and failure of the alert service, with staff expected to recognize the problem and use ordinary clinical escalation. The review records delays, over-reliance, workarounds, alert fatigue, and accessibility barriers. It also records failures to act when clinical judgment should take precedence. It establishes whether nurses, physicians, and support staff can disagree with an alert and whether the unit can respond at peak demand. **Figure 6** places these findings in the context of the model, interface, people, and working conditions.

The AI Steering Committee reviews the combined findings from the first two units and the proposed third unit. These include remaining risk and uncertainty, unresolved requirements, and every required specialist's signed or linked conclusion. The committee considers model and data-delivery performance together with how staff use the service, subgroup results, staffing and resources, supplier duties, the ordinary clinical alternative, and costs through retirement. It may approve only the units, patients, tasks, and scale supported by the evidence and within its delegated risk limits. A failed critical criterion, missing mandatory finding, or unmet condition due at this decision blocks approval of the affected use. The committee holds or rejects that part of the proposal, or asks for a narrower one.

A condition may remain outstanding only if it is legitimately due at a later stage. The record must state its owner, deadline, the activity it prevents, the evidence needed to resolve it, and who may accept that evidence. Every condition required to establish clinical safety or readiness for live use must be resolved before release. The governance decision and separate expansion authorization record the permitted addition, review or expiry date, and conditions. Neither allows the third unit to go live before production confirmation.

A1.5 Confirm readiness to expand

Before the third unit goes live, the internal technical owner checks that the prepared release matches the approved setup. Testing occurs in an authorized environment that reproduces live operation or in a controlled

release environment. It covers the exact model, input variables, threshold, eligibility and suppression rules, and scoring interval, as well as the interface, routing, identities, permissions, and monitoring calculations. Test messages must reach the correct roles across shifts without causing patient-care actions. Tests also show how the service handles missing scores and undelivered alerts. Users and operators must be able to tell whether data are current and whether the service is working, restricted, or unavailable. Logs must reconstruct the sequence from source values through scoring, suppression, delivery, and acknowledgement. Any difference between the test setup and release environment must be resolved, or assessed and accepted, before the test can support release.

The clinical and operational owners verify that forecast peak alert volume can be managed within the approved response standard. They exercise escalation and support coverage, downtime procedures, and the ordinary non-AI workflow, including reconciliation after restoration. The exercise must show that the hospital can return to the approved earlier configuration selected for rollback within the maximum recovery time. Staff must receive and understand the notice that the AI is unavailable and continue the ordinary deterioration pathway. A written procedure alone does not demonstrate either technical recovery or clinical continuity.

Staff demonstrate the abilities required for their roles before receiving access. Frontline users show that they understand the approved use, can assess or disagree with an alert, and can recognize and escalate concern when results are implausible or absent. They practise ordinary clinical escalation, downtime, and concern reporting. Technical staff reconstruct an alert and a missing-alert case. They also identify the live release, disable routing, and restore the approved baseline. Clinical and committee reviewers interpret threshold trade-offs, calibration, subgroup uncertainty, alert burden, and response capacity. The operational owner retains the assessment results and access decision for each eligible user. Completion depends on demonstrated competence, not attendance alone.

Monitoring feeds, incident reporting, support contacts, fallback, and pause authority must work before release, with recurring funding to sustain them. The budget covers model maintenance or hosting, integration and data engineering, clinical review, competency assessment, and ongoing support. It also covers incident response, rollback, continuity, and exit, including work still required when the service is restricted or retired. If funding ends, the hospital must reduce the service through an authorized, controlled change or retire it. It cannot continue after required monitoring or support has been withdrawn.

The three owners sign off on their responsibilities in **Appendix L Section 10.3**. The designated production confirmation authority then checks the release against both the governance decision and expansion authorization, using accepted evidence that every prerequisite for live use has been met. A missing or failed prerequisite results in Delay or Return for review. The confirmation records the exact release and setup, unit, users, permitted activities and volume, decision date, and activation date and time. It also records conditions due later, the dates of the first and subsequent operating reviews, and any expiry. The operational owner coordinates activation within these limits, and the dossier coordinator updates the registry and decision record. Production confirmation cannot extend an approval: another unit, population, threshold, or feature needs the applicable review and authorization before use.

A1.6 Monitor after expansion

The internal technical owner checks whether all eligible patients are being scored at the approved interval, comparing expected with completed scoring opportunities. Monitoring also covers missing or stale inputs, the version in use, routing, duplicates, acknowledgements, and receipt of all required monitoring data. Records must distinguish a valid below-threshold score, suppression allowed by the rules, and a missing alert caused by failure in scoring, suppression, delivery, or routing. Specified findings trigger investigation. Wrong-patient or wrong-team routing, an unverified version, or loss of a safety-critical data feed triggers immediate suspension and

incident escalation under the predefined rule. The operational owner tells staff which function is unavailable and confirms that they are using the ordinary clinical pathway. During an interruption, lack of an alert must not be presented as reassurance.

Clinical reviewers monitor calibration, sensitivity, positive predictive value, false-positive burden, and false negatives, alongside prevalence, useful lead time, and alert rate. Each measure uses its specified patient, encounter, event, alert, or observation-time denominator and reporting period, with repeated alerts and incomplete follow-up handled under the agreed rules. Reviewers obtain deterioration events from the defined outcome source, including events without alerts, and compare them with scoring and delivery records. They examine no-alert cases and unscored eligible patients separately from audits of received alerts. Calibration is assessed across the relevant scored risk ranges, not only among alerts.

Approved baselines, minimum-data rules, uncertainty rules, and subgroup results determine whether a conclusion is supported. A warning requires investigation and holds further expansion, while an action threshold requires its assigned mitigation. A serious-error or subgroup breach designated as a stop condition requires immediate suspension, fallback, and incident escalation. Missing or unreliable evidence invokes the approved response to uncertainty and loss of monitoring and cannot be reported as an absence of failure.

The operational and clinical owners review the interval from available source data to alert delivery, acknowledgement, bedside assessment, and indicated escalation. They compare each interval with its approved timing limit, separating technical delay from delay in clinical response. They check unresolved alerts and peak demand on every shift against response capacity. Audits include disagreements, overrides, manual escalation, workarounds, and deterioration without an alert, so they can detect under-reliance and over-reliance. Waiting for an alert, treating silence as stability, or bypassing ordinary escalation prompts case review and assessment of any need for further training or support. Persistent or severe failures may require restriction. This is particularly important when workload or workflow design prevents meaningful review and education alone cannot correct the problem.

The committee continues to assess value, outcomes, and equity against their own criteria even when technical performance remains acceptable. A missed value target may warrant redesign or retirement. Subgroup findings are judged against the minimum acceptable results and uncertainty rules set before evaluation. If event counts cannot support a conclusion, reviewers record the limitation and the further evidence or review needed. They do not infer equivalence from a small sample. They consider clinical results alongside patient and staff complaints, downstream tests, work passed to other teams, and effects on access.

The monitoring record must tell reviewers what to do with each result. For every measure it specifies approved warning, action, stop, and restart rules; primary and backup reviewers; the longest acceptable reporting delay; review frequency; and the response deadline. It names the escalation route and authorized decision-maker and states what evidence would justify operating under a restriction or restarting. The reviewer records the finding, decision, person responsible for the response, and follow-up until the action is accepted as complete.

Automated feeds may reconcile releases, detect missing data, calculate measures, notify owners, and link reports to the dossier. Authorized people judge clinical significance and decide the response within their authority. A stop condition requires immediate suspension, fallback, and escalation even before the cause is known. Any emergency continuity exception needs a separate documented decision by a person with the required emergency authority, must provide equivalent protection and an expiry, and cannot waive mandatory duties. The Tier 3 alert remains ineligible for streamlined lower-risk review.

An increase in alert volume does not by itself show that the model has deteriorated. The owners first investigate changes in deterioration prevalence, eligibility, codes or unit conversions, missing data, thresholds, platform versions, and user behaviour. They also investigate unexpectedly few alerts, missing scores, and deterioration

without alerts. The internal technical owner compares the live setup with the approved one and, where appropriate, tests the current and rollback versions against a controlled reference set and recent local data. Clinical and operational owners examine patient mix, practice, staffing, and responses. The supplier investigates the parts it controls within its agreed response time. This combined investigation helps establish whether the problem lies in the model, data, technology, or clinical workflow.

Findings determine the corrective work: a mapping defect requires repair and reconciliation of the affected period, while a supplier update may need to be deferred or rolled back as reviewers examine repeat tests of previously validated functions. A workflow problem may need staffing or interface changes, renewed competency assessment, or narrower use. Threshold changes, model recalibration or retraining, new predictors, changed suppression rules, new populations, and altered response protocols are material changes.

For each material change, the owners identify which assumptions behind approval may no longer hold. They obtain the required specialist review and assess the change's effects, including repeat tests of existing functions and subgroup evaluation. Tests and acceptance criteria are set before the new results are known. Release requires governance authorization and production confirmation, an identified earlier setup to restore if needed, and closer review afterward. The owners must resolve identified problems through correction, restriction, model or workflow revision, rollback, or retirement. If waiting to make an urgent protective change would create greater harm, the hospital may use the separately authorized emergency route. It must record the change immediately and complete follow-up review; mandatory duties still apply.

At the first operating review after expansion, the committee compares the third unit's evidence with its pre-specified criteria. These cover safety, performance, equity, user practice, workload, value, and support capacity. It records whether the unit may continue as approved, needs remediation or restriction, or should return to the ordinary pathway. Each required action has an owner, due date, and accepting authority. Enthusiasm or use volume cannot substitute for evidence, and no further unit enters scope without a later expansion decision and production confirmation.

A1.7 Manage incidents, restart, and retirement

In an illustrative incident, a source-system code change causes an oxygen-related field to arrive as missing, and the completeness monitor detects it. Under the approved stop rule, the internal technical owner immediately suppresses the affected function, preserves the evidence, and escalates the incident. Suppression prevents staff from receiving scores based on an unapproved data state. The operational owner notifies affected teams, makes the unavailable status visible, and confirms that ordinary observation and escalation continue with adequate coverage. The clinical owner directs review of affected patients, including those who deteriorated without an alert, to determine whether records need correction or clinical follow-up is required. The investigation establishes the start and end of the affected period and retains the model, source data, release, mappings, routing, user activity, and communications. If the hospital cannot establish which units or patients are unaffected, it widens containment until that scope is established. Applying the stop rule does not wait for confirmation of the cause.

The **Appendix H** incident record identifies the people and services affected, the severity of actual or possible harm, how quickly harm could occur, immediate protection, and required reports or notifications. It assigns hospital and supplier duties, including who may restrict the service or approve restart, and gives response deadlines and backup contacts. Communications and corrective actions remain linked to the investigation evidence. Investigators reconstruct the full sequence, examining input mappings, model behaviour, release history, workflow, and detection controls rather than accepting the first plausible explanation. They distinguish alerts never generated from those suppressed, delayed, misrouted, or not acted on. Each corrective action has an owner, due date, objective completion evidence, and a verifier independent of the work.

Restart can be considered only after the corrected input mapping passes tests using representative and affected cases, including cases with no alert, and after routing and monitoring tests also pass. The hospital must reconcile affected records, restore capacity to respond, and verify both the ordinary clinical alternative and the ability to stop the service again. A reviewer independent of the corrective work verifies the evidence. The authorized hospital decision-maker then accepts the evidence and records where and how restart is permitted, its conditions, review date, and any expiry. Severe incidents, material corrections, or changes in remaining risk return to the AI Steering Committee, which may narrow restart or require earlier review. Production confirmation for the affected version and use must be complete before alerts resume. Staff receive notice of restart, and closer monitoring checks that the correction works in practice.

The hospital considers retirement when benefit is not sustained, safety or subgroup criteria cannot be met, monitoring or response capacity is lost, a critical dependency cannot be governed, or consolidation or replacement is preferable. The authorized decision-maker records the retirement decision, scope, date, and owners of the closure work. The clinical and operational owners must verify that staff can recognize and escalate deterioration through the replacement or ordinary pathway and communicate the change to affected teams and patients where required. A stop condition requires immediate containment and cannot wait for retirement to be completed. To close the service, the owners disable scoring, interfaces, routing, identities, and access. They reconcile open alerts and affected records and complete data, log, contractual, and supplier duties. A named verifier confirms closure, including that no unsupported installation or running copy remains, before the registry records completed retirement.

Lessons from operation and incidents inform changes to procurement, monitoring, workflow, competence requirements, or policy. Education is refreshed after material changes, serious incidents, expansion, or observed practice problems. An authorized, de-identified lesson may be shared through **Appendix H** to inform later governance and sourcing decisions. The dossier coordinator keeps these findings in the service's continuing history and updates the registry's current authorized status after each decision. For an approved expansion, the record identifies the third unit and its configuration group, version, threshold, workflow, permitted scale, monitoring, and first review. It retains conditions, recurring funding, accountable owners, and the next extension that still needs approval. Later restriction, restart, change, or retirement decisions update the same record without erasing earlier authorizations. **Appendix I** defines the fields; **Appendix L** keeps the governance, production-confirmation, and expansion decisions distinct.

A2 AI scribe for clinical notes

In the second example, an adult ambulatory medicine service uses a purchased ambient scribe during approved in-person follow-up visits. The clinician deliberately starts recording, and the service uses the encounter audio to generate an unsigned draft in the correct EHR encounter. It is a generative AI use because it produces note text from the recorded conversation. The clinician must review and correct the entire draft before signing it. The service cannot finalize the note, populate structured fields, suggest codes, retrieve broadly from the chart, or create orders, referrals, tasks, or patient messages.

The hospital proposes extending the scribe to an adult cardiology clinic. This example assumes that after-hours documentation work and delayed note closure have been documented there; the committee must examine those baseline findings and the existing service's operating evidence. It must decide whether the proposed software version, settings, data connections, safeguards, and workflow can meet cardiology's needs, and whether staff can use and monitor the service safely. Supplier arrangements and recurring resources must also support the additional use. The review follows **Table 1**'s decision stages and **Figure 5**'s recurring tasks. [51][52]

A2.1 Define the proposed use

The hospital records the cardiology expansion request in the **Appendix L** record linked to the existing registry entry. The entry continues to show the current service's operational status; requesting expansion does not authorize cardiology use. Reviewers use **Appendix I**'s dossier fields to check the approved patients, tasks, and settings, the software and technical setup, data flows, and owners. They also check the approved maximum number of configuration groups the hospital can support, monitoring arrangements, the contract, and previous decisions. If current practice or enabled features already exceed approval, the hospital restricts that unapproved activity and brings the existing service back into line before considering expansion.

The record states exactly what the scribe may capture and produce, using the ambient-documentation choices in **Section 10.6.2**. In this example, a clinician must start recording; continuous recording is excluded. Reviewers establish whether dictated or uploaded content is accepted, which selected chart information may be supplied, if any, and how broad chart retrieval is disabled. They record audio and transcript retention and the supported speakers and languages. Only assessed input choices are approved. The sole permitted output is an unsigned draft, and the hospital verifies that structured-field population, coding, messaging, orders, referrals, and tasks remain disabled. Enabling any of these functions changes what the service can do and requires review before use.

The cardiology configuration record has its own version and identifies the supplier release, capture method, specialty template and prompt, supported language and devices, EHR interface, information supplied as context, and retention and support arrangements. It incorporates the input and output limits above, including unsigned drafts as the only output and disabled downstream functions. This becomes the reference against which reviewers assess changes to the live setup or what the service can do; such changes are not automatically ordinary maintenance. As **Figure 1** illustrates, one scribe platform may support uses that require separate approval and different risk tiers.

At Tier 0, the intake authority records whether the proposed review may proceed, must be held, or must be rejected, stating the reason, scope, corrective action, and who has authority to reconsider the decision. The proposal needs defined owners, a clear data path, and an assessable purpose; outstanding specialist work is referred to the responsible authority. A Proceed decision permits only the specified review or evaluation. Tests using patient data or connecting to live systems require the authority and safeguards applicable to that activity before they begin.

Before live expansion, the hospital must establish that production data processing is authorized and that evidence and logs are sufficient for oversight. It must verify that every draft remains linked to the correct patient and encounter and that controls enforce the required review before signature. The hospital must also verify ordinary documentation as the fallback, the ability to disable EHR insertion, and supplier support for changes and incidents when the relevant decision requires those controls. An unmet prerequisite holds the activity that depends on it. A prohibited use or a design that cannot be remedied is rejected.

Reviewers first consider Tier 4. They can exclude it only after verifying that the scribe cannot finalize a record or initiate a consequential action and that clinicians can reject a draft before anyone relies on it in a consequential way. Clinicians need the source information, competence, time, authority, and a workable alternative to do so. Tier 3 then applies because the signed note becomes an official record that can affect care, communication, coding, and follow-up. Important omissions, unsupported statements, wrong attribution, or insertion into the wrong patient's record can therefore cause harm beyond Tier 2. If a required limit has not been verified, reviewers retain the highest provisional tier supported by the credible route to harm and hold activities that depend on that limit. Automatic signature, merely formal review, or enabling structured-field writing, patient messaging, coding, referrals, orders, or tasks would reopen tiering and may create a separate action-taking use.

The clinical owner is responsible for the documentation standard and the professional limits on how the scribe may be used. The internal technical owner is responsible for the deployed configuration and data path. The operational owner is responsible for clinic readiness, support, fallback, and sustained capacity. The AI Steering Committee decides whether the proposed expansion is within its delegated tolerance, as shown in **Figure 2**. Privacy, cybersecurity, records, safety, professional, accessibility, equity, legal, and procurement functions retain authority over decisions within their own mandates. The committee cannot waive a required specialist finding. A decision beyond its delegated authority is referred to the authority identified in the hospital's decision matrix.

A2.2 Assess the need and possible harms

The cardiology service first establishes the documentation problem over a defined baseline period. It measures after-hours work, time to close notes, incomplete or corrected notes, clinician workload, downstream requests for clarification, and patient experience. Encounter volume and complexity help explain these results, as does observation of busy clinics, handoffs, and staff absences. The owner compares the proposed scribe with feasible alternatives: ordinary documentation or dictation, better templates or workflow, staffing support, or limiting use to narrower encounter types. The owner explains why the proposed setup is expected to address the measured problem more effectively. Access to an enterprise licence and clinician enthusiasm alone do not establish need.

Before reviewing results, the owners agree what improvement in after-hours work and note-closure time would justify expansion. They assess those outcomes alongside note quality, clinically meaningful corrections, and the time clinicians spend reviewing drafts. The evaluation also includes patient and clinician experience, effects on learners, work transferred to other staff, support burden, and full lifecycle cost. Uptake and editing rates help describe use but need interpretation. Low editing may reflect accurate drafts or superficial review. Reviewers therefore examine note quality and how clinicians actually verify the text before drawing a conclusion about value.

The hazard analysis follows an error from recording through the draft and signed note to the point where another person relies on it. It examines omissions or unsupported content involving medication, allergies, assessment, plans, and follow-up, as well as errors in whether a statement is affirmed or denied, who said it, and when events occurred. It also examines insertion into the wrong patient's record or encounter and recording outside an

approved encounter. Reviewers follow information through recording, draft, final note, and later amendment to identify where patient, encounter, or speaker identity could be lost.

Reviewers follow each error far enough to understand its possible consequences. They consider recording without first giving those present the required information, time pressure that makes clinician review a formality, downtime, and later reliance on an inaccurate signed note, including after correction. For each important route to harm, they identify how to prevent and detect the problem, the warning or stop condition, the responsible owner, and the response. **Figure 6** shows why the service must be assessed together with its interface, clinician, supporting team, and ordinary documentation alternative.

Patient and family partners review the information offered before recording and whether patients can realistically decline. Cardiology staff, interpreters, and learners contribute alongside health information professionals, coders, and downstream clinicians, who may see burdens or failure pathways that model testing misses. The value and hazard records retain their disagreements and unanswered questions so that the committee can consider consequential concerns when deciding whether to proceed.

A2.3 Assign hospital and supplier responsibilities

The example assumes that the supplier controls both the hosted core service and the models it relies on from other providers. The hospital controls the EHR connector, identity and access settings, approved template, and settings that enable or disable features. The hospital is also responsible for patient communication and clinical workflow. **Figure 3**'s component matrix records these separate responsibilities so that purchasing the service does not leave uncertainty about who controls each part.

The contract assigns the supplier responsibility for identifying releases, advance notice of material changes, and the model and service evidence needed for review. It also covers hosting, other providers that process data on the supplier's behalf (subprocessors), incident cooperation, defect correction, support for rollback or safely disabling the service, data return or deletion, and transition assistance. The hospital chooses and approves local use, validates the process used in live operation, controls settings and access, and monitors practice and outcomes. It retains decisions to expand, restrict, pause, restart, or leave the service. If a partner co-develops a specialty template or evaluation method, that component is recorded as a partnership. Each change, evidence item, corrective action, and handoff has one accountable party.

If a vendor proof of concept (POC) informed the purchase, reviewers link its authorization, protocol, results, amendments, and closure record through **Appendix G** and **Appendix L Section 5.2**. They check whether the tested setup and findings still apply to cardiology, including clinically important errors, supplier interventions, and departures from the protocol. Procurement evidence must be distinguished from the additional evidence needed to approve cardiology use. Any criteria changed after results were known remain disclosed amendments, and supplier-supported work is not described as independent evaluation. Reviewers also verify that temporary access and connections were closed and data were handled under the approved close-out plan. The procurement decision itself did not authorize patient-data access, clinical operation, or later expansion.

If no separate POC took place, the procurement owner and responsible governance authority record the evidence actually used and the reasons for omitting it under local policy. They do not describe that work as a retrospective POC. If a purchase or configuration commitment remains undecided, reviewers ask whether existing evidence or an equivalent authorized evaluation can resolve important questions about suitability, performance, integration, or workflow. If it cannot, a separate, time-limited POC is required.

The data-flow review follows audio, transcripts, prompts, drafts, and final notes separately from service-operation data, support logs, and derived data, because different access and retention rules may apply. For each type of information, reviewers establish where and in which jurisdiction it is hosted, which subprocessors handle it, who can access it, and how it is retained and deleted. They also establish whether the supplier may use it for

training or product improvement; secondary use stays disabled unless separately authorized. The technical owner checks that actual live data flows match the approved map and that incident evidence can be obtained within the required time. The records function confirms retention of signed notes and later amendments, keeping the official record distinct from audio and working drafts.

The expansion budget covers the ongoing work needed to sustain cardiology use, including licences, integration, independent note audit, training, and competency assessment. It also funds accessible patient information, support, incident response, continuity, supplier management, and exit at the authorized scale. If recurring funding ends, the hospital must narrow the service to a scope it can support or retire it, even if staff can still access the software.

This example assumes that the hospital has funded enough staff and resources for two configuration groups: the existing adult ambulatory service and cardiology. A separate cardiology group is justified only if its template, terminology, assurance, competence, monitoring, and recovery requirements differ materially. Within either group, evidence must show that variants can be evaluated, monitored, and supported together within approved limits. Unregistered prompts, unsupported languages, and new recording methods, information sources, or outputs remain prohibited until the required review and authorization are complete.

If the hospital reaches that two-group limit, it must consolidate or retire a group, fund and approve more capacity, or decline the proposed variation. An enterprise licence does not remove these staffing and resource requirements. **Appendix E** provides the lifecycle check; **Appendices I and L** record the maximum and current numbers of groups, why variants can share a group, and the decision. Funding two groups is a local departure from the framework's default planning limit of one. It is an assumption for this example, not a recommended limit for every scribe service.

A2.4 Evidence review and governance decision

Reviewers compare published and supplier evidence with the exact release, capture method, language, specialty, encounter type, workflow, and output proposed for cardiology. Evidence that a scribe reduces documentation burden can support the expected benefit, but it does not establish that the cardiology template is accurate or that drafts reliably enter the correct EHR encounter. Before approving expansion, the hospital also needs evidence that clinicians can detect serious errors during ordinary work. [52][53]

Before seeing new results, the clinical, technical, and operational owners agree how to test the cardiology service. The plan identifies the exact setup, sources of cases, sample-size justification, and cases that represent both ordinary and difficult cardiology encounters. It covers clinicians, encounter complexity, terminology, medication and follow-up content, speaker combinations, interpreters, accessibility needs, devices, rooms, and sound conditions. Unsupported languages and encounter types remain excluded. Reviewers need a reliable account of what happened in each encounter, obtained through standardized encounters, authorized recordings, or another suitable method. Collection must be lawful, and information is retained only as necessary for that purpose. Before testing begins, the plan defines error severity, acceptance criteria, and when results count as Pass, Fail, or Insufficient evidence. It separately defines conditions that require evaluation or live operation to stop.

Reviewers compare what occurred in the encounter with the generated draft, the clinician's corrections, and the signed note. They classify errors by clinical consequence rather than by grammar or the amount of textual change. The assessment covers important omissions and unsupported statements. It also covers errors in who said something, whether it was affirmed or denied, when events occurred, medication and allergy content, assessment, plans, and follow-up. The test cases must allow reviewers to check, for example, whether a denied symptom becomes an asserted symptom, a family member's history becomes the patient's history, or a medication change or follow-up instruction is lost.

The report states the denominator used for each result, such as the number of encounters assessed, and the uncertainty in that result. It reports any deviations from the plan and conditions not tested, with separate specialty or language results where required. An acceptable overall result cannot compensate for a serious-error category or subgroup result that fails its pre-specified criterion. Insufficient evidence cannot be treated as a pass. If the gap prevents a required finding or leaves an applicable critical criterion unmet, the affected decision or use remains on hold.

Reviewers examine whether the new notes affect coding, billing, quality reporting, or compliance, and whether other staff must seek more clarification than before. Health-information, coding, or compliance specialists record findings wherever their functions are affected. The records team and responsible clinicians also test correction of an inaccurate signed note. They verify that the original entry and amendment history remain traceable and that the correction reaches people who need it for care or other work. A readable note is not sufficient if it introduces errors into the official record or creates substantial correction work for other teams.

Representative clinicians test the proposed release in an authorized environment that includes routine and complex encounters, interruptions, time pressure, disagreement, errors, and downtime. Reviewers observe whether clinicians recognize draft status, can obtain the source information needed to verify a note, and can correct or reject fluent but inaccurate text. Tests follow patient and encounter identity from recording through insertion, review, and signature, including deliberate wrong-patient or wrong-encounter mismatches. They assess both interface safeguards and clinician performance. Reviewers then establish whether the expected workload leaves enough time for proper review. They examine workarounds, excessive trust in drafts, effects on learners' independent documentation skills, and correction work passed to other staff. Together, these findings help establish whether the proposed review process can work in the clinic.

Privacy and professional-practice reviewers establish the applicable legal basis and whether the hospital requires encounter-specific notice, a process for objection or opt-out, explicit consent, or another authorized approach. Before recording, patients should be able to understand what is captured and produced, how their information is handled, who checks the draft and remains responsible, and how to decline, ask questions, or request correction. The workflow accounts for interpreters, family members, other speakers, and people present who are not participating, and provides accessible non-AI documentation without disadvantage. [51]

Where feasible, approved information is available before the appointment so that making this choice does not displace the clinical purpose of the visit. The clinic also establishes how it will handle a late objection, interrupted recording, or correction request without disadvantaging the patient. Staff must complete the required information and choice process before capture begins. Agreement obtained later does not show that the required process was completed before recording.

The committee may authorize cardiology use only if the evidence and specialist findings support it and the remaining risk falls within its delegated authority. The decision specifies the approved setup and limits on where, by whom, and for which encounters it may be used, its Tier 3 classification, the cardiology configuration group defined above, and the features that must remain disabled. It records evidence limitations, remaining risk, stop rules, the first review date, and any expiry. Each condition has an owner, due date, the activity it prevents, and the authority that will accept evidence of completion. If a condition for live use remains unresolved, only expressly permitted preparatory work may continue. The clinic cannot go live until every release prerequisite is met and verified and separate production confirmation is complete.

If expansion is authorized, the initial group of cardiology clinicians must have a demonstrated documentation burden and work representative of the proposed service. Its leaders must be able to maintain the approved workflow, with enough staff and support to contain a problem. The group and its operating limits are selected before production confirmation and recorded there. This defines who may begin using the service, for which encounters and at what scale, with a planned review. Enthusiasm alone does not establish readiness.

A2.5 Confirm readiness to expand

Production confirmation checks the service prepared for cardiology against the approval. The owners verify the supplier release, models, prompt, specialty template, devices, users, EHR interface, patient and encounter links, retention arrangements, and feature settings. Testing takes place in an authorized environment representative of live use; any differences must be assessed and accepted before the results can support release. Tests must show that each output enters only the correct encounter, is visibly an unsigned draft, and cannot bypass required clinician review and signature. Coding, structured fields, messages, orders, referrals, tasks, chart-wide retrieval, and automatic signature remain unavailable unless separately governed. Testing must not enable unapproved patient-facing activity. Any mismatch with approval must be corrected, referred for focused review, or held from release.

The owners complete **Appendix F**'s scribe production checks and link the results to the cardiology decision. They verify patient information before recording, visible start and stop controls, arrangements for families and interpreters, correction and complaint routes, and access to ordinary documentation. They confirm that the test sample and error-severity rules apply to the setup being released and that results for important errors meet the pre-set criterion. Duties for audit, supplier notice, rollback or safe disabling, record correction, restart, and exit must have named owners and response times before confirmation is complete. Release evidence must include successful tests of disabling insertion, continuing ordinary documentation, and completing the approved record-amendment procedure.

Monitoring must work before staff receive access to the live service. The hospital verifies counts of eligible encounters and encounters with signed notes, establishes an independent note-audit sample, and checks that it receives information on links between drafts, patients, and encounters, feature and release changes, patient questions or objections, downtime, support demand, and the value measures chosen at intake. For each measure, the **Appendix I** record defines the population, denominator, source, baseline, and uncertainty, explains how completeness is checked, and sets the longest acceptable reporting delay and review interval. It specifies warning, action, and stop responses and names reviewers and the people authorized to act. The hospital also tests detection of a missing critical data feed or audit log and the response to that loss. The approved stop rule applies even before the cause is known.

Staff receive access to the live service only after release is confirmed and they have demonstrated the skills their roles require. Training and testing use separately authorized access. A cardiology clinician must correct or reject a fluent draft with an important omission, wrong attribution, or wrong-encounter text, and complete ordinary documentation during an outage. Technical staff demonstrate how to identify the live supplier release, trace the authorized data route, reconstruct an insertion, preserve incident evidence, disable EHR insertion, and restore the approved setup. Operational leaders show that staffing covers audit, support, escalation, patient communication, and ordinary documentation at the proposed volume. Relevant competence is reassessed after a material change, serious incident, prolonged absence, or meaningful expansion.

The production record states which cardiology users, clinic, encounters, and language are covered, together with service hours, maximum volume, and the first review date. It confirms recurring funding and sufficient staff and resources for monitoring, audit, support, incident response, ordinary documentation, and rollback or safe disabling at that scale. The three owners sign off on their responsibilities. The designated production authority then records in **Appendix L Section 10.3** that the released service matches approval and that every condition due for live use has been met and verified. Only the recorded cardiology use may then begin.

A2.6 Monitor initial use and expansion

During an encounter, clinicians verify the patient and encounter, follow the approved patient-information process, and control when recording starts and ends. Before signing, they review the entire draft against the

encounter and the authorized information needed to verify it. They correct every material error and confirm the patient and encounter again. If a statement cannot be verified, they resolve it or remove the unsupported text before signing. A fluent draft is not evidence that an event occurred. They use ordinary documentation if the service is unavailable or the encounter is outside scope. They report recurring errors, privacy concerns, unexpected features, workarounds, and downstream problems through the designated route. Errors found after signature follow the authorized record-amendment and clinical follow-up process. Repeated use outside scope or superficial review leads to case review, coaching, restriction, or removal of access according to severity.

The interface supports this work by making recording status visible, enforcing draft status, and displaying patient and encounter identity during review. It records the clinician's review and sign-off so that the final note and any later amendment can be traced to the correct encounter. Access controls restrict use to approved users and settings. Targeted note and use audits examine unusually low editing, rapid signing, repeated corrections, complaints, and use outside approved encounter types. Educators also preserve opportunities for learners to practise independent documentation and clinical synthesis, because those skills remain necessary for verification and fallback.

Routine review asks whether notes remain safe, clinicians can still verify them properly, patients can exercise their choices, and the service continues to provide the benefit that justified expansion. Reviewers consider errors and their severity alongside review time, delayed signing, record corrections, and patient concerns. They use downtime, support demand, work transferred to other staff, and supplier releases to help explain changes in performance or burden. Where the data allow meaningful interpretation, they examine results separately by specialty, language, clinician cohort, and other locally relevant groups. The dossier records findings, required responses, and any evidence limitation that affects the decision to continue.

A warning, such as an increase in a recurring type of omission, requires investigation before further expansion. That investigation uses targeted sampling, clinician feedback, and template review. At an action threshold, the owner carries out the predefined response; depending on the measure, this may restrict a template, device, room, encounter type, or user group. A stop condition requires immediate suspension of the affected recording or EHR-insertion function, ordinary documentation, and incident escalation and response. Examples include wrong-patient insertion, bypassed draft status, unapproved recording or data flow, repeatable serious errors, and inability to reconstruct use from the available records. The record states what evidence and authority are needed to remove a restriction or approve restart. Restart requires accepted corrective evidence, a recorded decision by the authorized body, and completed production confirmation before the affected operation resumes.

At the first operating review, the committee decides whether cardiology use may continue, needs correction or restriction, may expand to the remainder of the approved specialty, or should end. Expansion requires the pre-specified results for safety, note quality, patient choice, use within approved limits, workload, equity, and value. Supplier support, monitoring resources, and recurring funding must also be sufficient. No unresolved stop condition or overdue corrective action may carry into expansion. **Appendix L Section 10.4** records the expansion decision separately and states which next extension still needs authorization. Any approved extension begins only after production confirmation for the added users, locations, volume, or other affected parts of the service.

A2.7 Review configuration changes

When performance changes, investigators examine supplier changes, hospital changes, and clinical practice. A supplier release may alter speech recognition, the language model, prompts, default templates, retrieval, or enabled features. Hospital changes may affect the EHR connector, patient and encounter links, draft controls, devices, or feature settings. Practice may also move beyond the validated specialty, language, room conditions, speaker mix, terminology, or review practice without any software update. Any of these can affect note quality,

privacy, or later reliance on the record. The approved record therefore links the supplier version and local setup to the conditions of use and observed review practice. Investigators use it to establish what changed and who must correct the problem.

The hospital keeps a versioned set of reference test cases and checks for the approved live setup. These allow it to assess supplier notices and proposed local changes. The initial assessment identifies which component and use are affected, which assumptions behind approval may no longer hold, what evidence and specialist review are needed, and the earlier setup available for rollback or the approved plan for safely disabling the service. It records whether focused review, renewed validation, or a new governance decision is required and who may decide.

Low-impact maintenance may remain within the existing configuration group if focused repeat testing shows that previously validated functions still work. A change that could affect note content, patient or encounter identity, clinician review, data use, or what later uses or actions are permitted requires targeted revalidation and the required change decision. Production confirmation must be completed before release. A new output or action is reviewed as a new or materially changed use, and the affected feature stays disabled until review is complete.

The party controlling the affected component investigates and corrects it. The supplier provides evidence and support for rollback or safe disabling when addressing its own defects. The hospital investigates its connector, template, settings that enable features, devices, data flow, and user practice. It may restrict use while the cause is uncertain. A stop condition still requires immediate suspension. For a co-developed component, the partnership matrix identifies who leads the investigation, how evidence is handed over, who owns the correction, and how disputes are handled. The hospital retains its authority to pause. The findings lead to the necessary correction, restriction, rollback, or revalidation before any required restart decision.

A2.8 Manage incidents, restart, and retirement

In an illustrative incident, an EHR connector update inserts a draft into the wrong encounter. The hospital immediately disables EHR insertion for the affected configuration group, returns to ordinary documentation, and records the pause and incident escalation. Before evidence is lost, the incident team preserves the supplier release, connector version, and logs of identities and draft-to-encounter links. It also preserves affected drafts, access records, and the change record. It identifies every potentially affected encounter, including whether an incorrect draft was signed or used downstream. The records team coordinates the authorized amendment process, and responsible clinicians correct the official record while preserving the history of the entry and correction. The clinical and operational owners arrange necessary clinical follow-up and patient communication. The responsible authorities complete any applicable privacy, safety, legal, and regulatory reporting.

Investigators examine the supplier service and EHR connector together with change review, production confirmation, clinician practice, monitoring, and approval conditions. Corrective work goes to the party responsible for the affected component. Before restart, the repaired service must pass tests of patient and encounter links, safeguards against wrong-patient insertion, and controls that keep excluded features disabled. Tests must also show that previously validated functions, ordinary documentation, and monitoring still work. A reviewer independent of the repair verifies objective evidence of completion. After the responsible authority accepts that evidence, the authorized body records whether restart is approved, where it is permitted, and its conditions. Production confirmation must be complete before EHR insertion resumes, followed by the specified closer monitoring. A successful repair or test alone does not authorize restart.

Appendix H retains the incident chronology, containment, affected-record findings, corrective actions, and restart decision, linked to the production confirmation and follow-up findings. The shared learning record explains what other hospitals may need to change in connector testing, supplier terms, monitoring, or literacy. A

de-identified account is shared through the authorized learning route without disclosing identifiable patient information.

The hospital considers retirement when benefit no longer justifies cost and burden, recurring funding or supplier support is lost, or a serious problem cannot be corrected. Loss of required monitoring or data rights, or consolidation of configuration groups, can also justify closure. The decision-maker authorized under the hospital's decision matrix records the retirement decision with operational leadership, including transition, communication, and closure duties.

The owners then disable access, settings that enable features, EHR insertion, interfaces, and service identities. They verify that no unsupported installation or running copy of the hospital's configured service remains. They reconcile open notes and corrections, export required records, retain required information, verify required deletion by the supplier and subprocessors, and complete contractual duties. Clinicians must be able to use the replacement documentation process, while affected staff and patients must receive the required transition communications and follow-up. The registry records completed retirement only after closure evidence is verified, and the dossier retains lessons learned. Stopping new enrolment alone does not fulfil these obligations.

The dossier must allow a later reviewer to identify exactly which scribe version, settings, data connections, and workflow are authorized, why the use is Tier 3, and what evidence supported cardiology expansion. It must identify who is authorized to act when conditions change and which functions remain unapproved. **Appendix L** connects this working record to **Appendix I**'s common fields, **Appendix F**'s scribe production checks, and the other applicable evidence and lifecycle records.

A3 AI agent for post-discharge coordination

In this example, an adult transitions-of-care service uses an AI agent to reduce missed or delayed follow-up after discharge. Unlike a system that only supplies information, the agent can coordinate steps and act through scheduling and communication tools, within the limits described here. It starts only after a clinician signs a discharge plan and creates a structured follow-up task. It may read the patient and discharge-episode identifiers, signed instruction, approved scheduling information, communication preference, and relevant referral status. It may check completion, retrieve available appointments, and prepare a standard message. Before any consequential action, a coordinator checks the source instruction, patient, episode, destination, appointment, and message. The coordinator separately authorizes booking and communication. An appointment-confirmation message may be sent only after the booking has been verified in the scheduling system. If the booking or message changes, the affected action requires fresh review.

The service is assumed to be approved for one transitions team, and the hospital proposes extending the same actions to a second team. The request adds no tools or data, does not allow free-form messages, and does not reduce human confirmation. Other functions available on the supplier's platform remain unapproved. **Table 1**'s seven decision stages and **Figure 5**'s six recurring tasks apply to this defined service and its expansion. [54][55][56]

A3.1 Define actions and assign the tier

Reviewers reopen the existing registry entry and dossier rather than creating a separate account of the service. **Appendix L** records what the first team is using, the proposed second team, requested volume, and exact expansion under review. The goal, input fields, tools, and permitted actions remain unchanged. So do the requirements for coordinator approval, permissions, information retained across steps or tasks (memory), transaction limits, handling of exceptions, and manual alternatives. **Appendix I**'s common fields let reviewers compare the second team's proposal with the existing approval. The record also specifies which steps must occur before others, what evidence proves each step is complete, and when unresolved work must pass to a person.

The agent's approved functions exclude clinical interpretation: it cannot interpret symptoms, give advice, determine urgency, change clinical priority, or order medication or tests. Nor can it create or cancel referrals, search unrelated records, send free-form messages, add tools, or act on unsigned instructions, even if the platform offers those capabilities.

Work the agent cannot complete within its approved limits goes to a staffed exception queue: a list of tasks requiring a person's attention. This includes replies about symptoms, conflicting instructions, unavailable destinations, identity mismatches, and unresolved ambiguity. A named staff member must accept the handoff; adding a task to a queue does not prove that anyone has received or acted on it. Symptoms and clinical questions go to the staffed clinical service under the clinical owner's escalation procedure, without the agent deciding urgency. Each transaction ends with a recorded status: completed and verified, rejected for manual handling, escalated, or expired. These describe the agent's work, not whether the patient received the needed follow-up. Rejected, escalated, and expired tasks keep an accountable human owner until the underlying work is resolved. The agent cannot retry indefinitely.

At Tier 0, reviewers check that the proposed review and evaluation have named owners, a clear purpose, authority to use data, and enough information for assessment. Before any activity begins, they verify the permission limits, records, independent means of stopping the agent, and other controls needed for that activity. Before live expansion, they must additionally verify limited permissions, records that reconstruct attempted and completed actions, and the hospital's ability to block further changes to other systems without relying on the

agent. Reviewers must also confirm supplier or partner incident support and enough capacity to handle the forecast workload manually.

The Tier 0 record states Reject, Hold, or Proceed, the activity and conditions covered, the reason, any corrective action, and who may reconsider the decision. Proceed permits only the specified review or evaluation. An unmet requirement holds the dependent activity until the requirement is met or the proposal is narrowed, redesigned, or covered by a revised contract. A prohibited use or a design that cannot be remedied is rejected.

Reviewers first consider Tier 4. They can exclude it only after verifying that an authorized coordinator makes a separate decision for every consequential booking and message. The interface must show the signed source instruction beside the complete proposed action, and the coordinator must have the time and authority to reject it. Technical controls prevent bypassing or reusing approval, while urgent tasks and tasks requiring clinical interpretation remain excluded. Reviewers must reconsider Tier 4 if the agent determines urgency, makes clinically significant decisions about access, or acts before meaningful review. They must do the same if batch confirmation or inadequate information makes approval a formality. The evidence and approval for this restricted Tier 3 service cannot be carried over to that different use.

Tier 3 applies because a completed booking or patient message can materially affect access, follow-up, and clinical communication. Those consequences exceed Tier 1 or Tier 2 even when verified safeguards reduce residual risk. **Appendix L Section 4** records the rationale and remaining uncertainty, explains where the service would cross into Tier 4, and identifies changes that require re-tiering.

The clinical owner sets task eligibility and clinical escalation requirements. The internal technical owner is responsible for the rules governing agent actions, permissions, integrations, logs, and monitoring. This includes maintaining the ability to prevent the agent from changing other systems (disabling writes) and restoring an earlier approved setup if needed (rollback). The operational owner maintains the coordinator and exception-handling services, assigns named staff for coverage, ensures that handoffs are acknowledged, and responds to overdue work. The AI Steering Committee decides whether the proposed actions and remaining risk are acceptable within its delegated authority. Scheduling, privacy, cybersecurity, safety, equity, accessibility, records, legal, and procurement specialists retain authority within their own mandates, as **Figure 2** shows.

A3.2 Assess the need and possible harms

The clinical and operational owners first establish how often the current pathway misses the prescribed follow-up interval and how long referral and scheduling tasks remain unresolved. They examine the manual contacts required, where work accumulates, and which patients experience the longest delays. The second team's baseline is recorded separately. Its population, discharge patterns, available appointments, languages, and workload from tasks that require manual handling may differ from those of the first team.

The hospital compares the agent with feasible improvements to the current process, such as clearer discharge tasks, more coordination staff, centralized scheduling, reminders, or better integration without action-taking AI. Before results are reviewed, it defines measures, denominators, the observation period, targets, and the response to a missed target. The primary outcome is follow-up completed within the clinically specified interval. The clinical and operational owners assign responsibility for verifying completion and responding through the established clinical pathway when follow-up is missed or unconfirmed. A booking or sent message alone cannot establish completion. Reviewers also assess delays, lost follow-up, no-shows, manual contacts, coordinator time, and waiting time for unresolved tasks. They consider patient experience, corrections, full lifecycle cost, and work passed to clinicians, schedulers, interpreters, and support staff. Processing more tasks alone does not establish benefit.

The hazard record follows failures in identity, destination, timing, content, and transaction state through to their possible consequences. A failure may involve booking for the wrong patient, discharge episode, or service, or

sending an inaccurate message or one addressed to the wrong recipient. The agent may miss a conflict, omit a required step, or carry out dependent steps in the wrong order. An interrupted update to another system may be reported as complete, an exception may be left unattended, or a retry may create a duplicate. The record also addresses stale or incorrect task memory, messages sent but not delivered, and handoffs or required patient responses that never arrive.

For each important route to harm, the record identifies how to prevent and detect the problem, who must respond, and how much time is available. The clinical owner sets clinical response requirements; the operational owner assigns staff coverage and escalation for missed acknowledgements and overdue work. The agent does not set a follow-up interval or decide a clinical urgency threshold. **Appendix L** records the problem, permitted use, and evidence decision, with detailed evidence and lifecycle fields in **Appendices D–F and I**.

The owners consider equity when defining the problem, including whether the manual service already serves some patients poorly. Reviewers examine whether barriers involving portal access, language, disability, geography, appointment availability, or digital literacy would divert some patients into a slower queue for staff attention. The baseline therefore measures how long tasks wait for manual handling and what happens to them, as well as completed follow-up. This allows reviewers to detect whether faster simple transactions conceal longer delays for other patients.

A3.3 Assign service responsibilities

In this example, the supplier provides the hosted model and the software that coordinates the agent's steps, called orchestration. The hospital builds the local extract of follow-up tasks, connector, and exception queue and maintains links to the correct patient and discharge episode. It controls monitoring and an independent way to stop the agent from changing other systems. A regional partner maintains the agreed rules for permitted actions and the transaction test library. **Figure 3**'s component matrix assigns these duties separately, so the hospital can identify who is responsible for each part.

For every material component, the responsibility matrix names both the party doing the work and the accountable hospital owner. The supplier is responsible for its hosted service, releases, upstream providers, defect correction, and the technical evidence and incident support required by contract. The regional partner maintains the versioned action rules and agreed tests. The hospital controls credentials, integrations, confirmation design, patient communication, operating limits, and independent monitoring. It also controls the allowlist: the list of tools and actions the agent is permitted to use. Decisions to release, expand, restrict, restart, or retire the local service remain with the hospital.

The internal technical owner coordinates investigation even if another party controls the suspected fault. Agreements must let the hospital identify each release, receive notice of material changes, and obtain evidence about actions and the full configuration in use. They must specify response times by severity, correction and rollback support, continuity arrangements, and ways to resolve disputes or leave the service. The evidence must show what was proposed and approved and what the receiving system actually did, including delayed responses and actions still in progress. Intellectual-property or publication terms cannot prevent reconstruction of an action, providing evidence to an authorized investigator, continuity, or exit. Expansion stays on hold if an important duty has no responsible party or the hospital lacks enough evidence to exercise its authority.

The expansion budget includes the hosted service and integration maintenance, monitoring and log storage, transaction audit, and competency assessment. It also funds coordinators and the exception queue, supplier and partner management, incident response, accessible communication, continuity, and exit. For each recurring funding source, the record identifies its term and owner and the renewal date at which continued support must be decided. If support ends, the hospital must reduce action volume or return work to manual coordination. It cannot continue with insufficient monitoring or queue capacity.

This example assumes that the hospital has approved enough staff and resources to support one configuration group for the agent's coordinating software and actions. Rosters, queue assignments, and hours may vary only within verified operating limits. The group's approved setup specifies the model, action rules, tools, required data formats and fields, messages, permissions, memory arrangements, confirmation design, and transaction limits. A new connector, action, message, permission range, or memory arrangement can remain in the group only if evidence shows that the hospital can evaluate, monitor, and support it with the existing variants. Otherwise it needs a separate group. At the approved maximum, the hospital must consolidate, retire, or decline the variation, or fund more resources and approve a higher limit. A temporary exception cannot waive sufficient staffing and resources, core controls, or required authority. As in Case 1, one group is an assumed local planning limit, not an empirically established maximum.

A3.4 Test actions and decide what may proceed

Before expansion, reviewers compare the second team's patients, discharge tasks, and clinical urgency with the approved service. They also compare languages, accessibility needs, appointment destinations, connector behaviour, hours, volume, and work requiring manual handling. Evidence from the first team must show use within approved actions, meaningful coordinator review, and the required responses to threshold breaches. Serious incidents and corrective actions must be closed, and results must support timely, equitable follow-up and continuing value. The second team then demonstrates competent staff, sufficient clinical and exception coverage, tested manual capacity, monitoring, incident support, and a setup consistent with the approved configuration group. Local tests cover actual receiving services, acknowledgement and escalation, unavailable staff or destinations, and checks that bookings and messages agree with receiving-system records at the proposed volume.

Evaluation follows each complete transaction from the signed discharge task through information retrieval, the proposed plan, and coordinator review, then to the request sent through a tool, the receiving system's response, the recorded outcome, and checks that the records agree. Test scenarios include correct and conflicting instructions, unavailable or stale appointment slots, duplicate tasks, expired approvals, and delayed or missing tool responses. They also cover changes to scheduling-data structure or required fields, interrupted updates, partly completed transactions, omitted steps, and steps attempted before their prerequisites are complete.

An apparent success must be checked against the receiving system's actual record for the correct patient, discharge episode, destination, and appointment. Acknowledgement that a tool received a request does not prove it completed the action. Tests distinguish a confirmed booking, a message accepted for sending, delivery verified where the channel supports it, a required patient response, and completed follow-up. If completion cannot be established, the outcome remains unresolved until a person checks and reconciles it. The agent must not guess or repeat a booking that may already have succeeded. Testing uses permissions equivalent to those proposed for live operation without creating real patient actions.

A separate assessment deliberately tries to make the agent break its operating limits. Testers place instructions in a discharge task, referral, or tool response that attempt to redirect it. They try to obtain another patient's information, use the wrong discharge episode, substitute tools, gain extra permissions, bypass required approval or reuse outdated approval, retry repeatedly, hide extra actions, or carry information between patients' tasks. They also test whether old task status or remembered instructions can trigger action after the signed plan changes or the task ends. Results distinguish correct completion from wrong targets, omissions, duplicates, incorrect action order, unauthorized tool calls, and unresolved partial transactions. Safe transfer to the exception queue is assessed separately and requires an identified receiving person or service to confirm acceptance of responsibility. Evaluators measure the coordinator time and workload associated with these outcomes.

Some failures block authorization regardless of average performance: bypassed approval, cross-patient exposure, use of a tool outside the allowlist, a wrong-patient external action, or a transaction that cannot be

recovered. Each defect must be corrected and retested. For other outcomes, the hospital sets locally justified limits on failed actions, exception-queue workload, recovery time, and reviewer errors. These limits and serious-error rules must be fixed before evaluation begins. They cover every applicable action type and specify when unresolved actions or unacknowledged handoffs must pass to a person, using the service's clinical response requirements. **Appendix F**'s action-taking and agentic checklist identifies the evidence needed for the action register, allowlist, confirmation, memory, reconstructable records, simulation, shutdown, and recovery controls.

Representative coordinators test whether they can review proposals meaningfully during normal, ambiguous, time-pressured, interrupted, and downtime work. They must identify the source instruction, recognize a wrong patient or destination, and reject unsupported proposals. They must also route symptoms appropriately, recover incomplete work, and continue manually when needed. Exercises include a receiving service that does not acknowledge a handoff and a patient who does not respond when a response is required. Staff demonstrate the established escalation or alternative-contact procedure and record who takes responsibility. Reviewers record and assess errors, delays, workarounds, accessibility barriers, rushed approvals, and exception workload against the agreed criteria. **Figure 6** places those findings in the context of the agent, interface, coordinator, clinical escalation route, and manual alternative working together.

Privacy and cybersecurity reviewers trace each transaction from the signed discharge task to the external booking or message and back to a verified outcome. The data map covers patient data, prompts, tool responses, messages, logs, and temporary information about the task and its progress. For each, it states how the information is processed, who can access it, and how it is retained and deleted. It defines what the agent may remember about a task, how that information stays linked to the correct patient and episode, and when it must be refreshed, cleared, or retained only in a protected audit record. Information from one patient or episode must not silently become input for another.

The agent receives only the access needed for the correct patient, episode, tool, and action, preferably through credentials that expire after a short period. It cannot browse generally, add tools, carry one patient's information into another patient's task, or keep broad ongoing permission to change other systems. Checks of required fields and data formats, together with alerts for unusual tool use, help identify departures from these limits. The hospital must be able to revoke credentials or block writes independently of the agent so it can contain a problem before knowing its cause. Tests must establish what happens to queued requests and actions already in progress when that control is used. Stopping the agent does not reverse actions already completed elsewhere.

The equity and accessibility assessment examines whether benefits reach patients beyond those who can readily use the portal or fit the simplest scheduling pathway. Reviewers examine results across locally relevant languages, disability and accessibility needs, communication channels, geography, and any other groups identified in the local review. They assess completed follow-up alongside the time tasks spend waiting for manual handling, and ensure that patients can use an accessible non-digital route without disadvantage.

Patients receive clear information explaining that automation contributed to the booking or message and that a named team remains responsible. It explains how to ask questions, correct errors, or use ordinary coordination, whether replies are monitored, staffed hours, and the existing route for urgent clinical help. An automated acknowledgement must not imply that a clinician assessed a symptom or accepted responsibility for it. If a required patient response does not arrive, staff follow the approved contact and escalation procedure. Silence cannot be recorded as agreement or completed follow-up and does not justify abandoning the task. Decisions about notice, consent, substitute decision-makers, records, privacy, accessibility, and Indigenous governance remain with the authorities identified in the hospital-specific governance arrangements described in **Appendix J**. Each decision requires approval from the authority responsible for it.

The committee reviews why the service is Tier 3, its expected value, specialist findings, transaction tests, and evidence that staff can review and respond reliably. It also considers remaining hazards, component

responsibilities, funding, and monitoring. It may approve the same actions for the second team only if the findings and remaining uncertainty support that use within its delegated authority. Before release, the hospital must show that it can handle the expected exceptions, complete local pre-release checks, and enforce the partner's duty to support rollback. The record names owners and the authorities that accept evidence of completion; every outstanding condition identifies its due date and the activity it prevents. The decision specifies the approved setup, permitted and excluded actions, review date, and the next extension that remains unapproved. Separate production confirmation is still required before the second team goes live.

The expansion decision identifies the second team, users, hours, volume, and permitted actions, together with conditions, the first operating review, and the next extension that still needs approval. The team may go live only after every release condition is met and verified and production confirmation is complete. If evidence or staffing and resources are insufficient, the committee holds or narrows expansion. The first team may continue only while it meets its own approval and operating conditions. Any problem discovered during expansion review must also be assessed for its effect on existing use.

A3.5 Confirm readiness to expand

The technical owner checks the second team's prepared release against the approved model, software that coordinates actions, action rules, and required data format and fields for each action. The comparison covers permitted tools, connectors, permissions, memory settings, message templates, operating limits, logs, and the earlier setup available for rollback. Testing then takes place in an authorized environment equivalent to live operation or in a controlled release environment. Tests verify patient and episode identity, approval of each individual action, duplicate prevention, and checking of the receiving system's record after each change. They also verify expiry of approvals, independent blocking of changes, and transfer to the exception queue. Testing covers action order, rejection of missing prerequisites, absent or delayed acknowledgements, and reconciliation of work interrupted when writes are disabled. None of these checks may create unapproved patient-facing actions.

Before release, monitoring, alert delivery, incident reporting, and supplier contacts must work. The manual queue, reconciliation arrangements, and a read-only mode that prevents changes to other systems must also be ready. The owners demonstrate recovery within the approved time and show that staff can continue outstanding tasks manually without losing patient context. Staff must distinguish actions confirmed complete, actions known to be in progress, actions not attempted, and actions with unknown outcomes. Before retrying or completing uncertain work manually, they check it against the receiving system to prevent omissions or duplicates. The exercise identifies who accepts each outstanding task and how an unacknowledged handoff is escalated. A failed control must be corrected and verified or returned for focused review. The affected release remains on hold until the requirement is met and verified, even if governance approval was already granted.

Before receiving live-service access, coordinators pass a scenario-based competency assessment. Training and testing use separately authorized access. Supervisors demonstrate how they will review rapid approvals, overrides, corrections, unresolved exceptions, and use outside scope. Technical operators demonstrate how to reconstruct transactions, revoke credentials, disable writes, roll back, and recover the service. Clinical owners show when tasks must leave the agentic pathway and how the clinical escalation service will be staffed. Patient-facing staff explain the service, alternatives, and correction route. These demonstrations establish readiness to perform the work. Attendance alone is insufficient.

For each team, **Appendix L Section 10.3** records the authorized actions, hours, volume, queue limit, support, first review, and authority to block writes. The second team requires its own production confirmation, recorded by the named confirmation authority after the three owners sign off. The expansion authorization in **Appendix L Section 10.4** does not itself permit release; the hospital must first verify that the second team's actual setup and working conditions meet its approval.

A3.6 Support users and monitor operation

The coordinator reviews one transaction at a time. The interface displays the signed instruction alongside the patient, episode, destination, proposed appointment, message, and any uncertainty. Technical controls prevent batch approval or reuse of approval, expire stale proposals, and require a fresh decision after a material edit. Approval covers only the action and details reviewed. A failed or unverified booking cannot trigger a confirmation message. Ambiguous work goes to manual handling; symptoms and clinical questions go to the staffed clinical queue. The coordinator remains responsible for the handoff until the receiving service accepts it and follows the established escalation procedure if acknowledgement is absent or overdue. Coordinators reconcile exceptions with the receiving system's authoritative record and use the manual pathway during downtime.

Supervisors examine actual review practice alongside competency records. They look at source access, changes or rejections, how long exceptions remain unresolved, and whether completed actions matched their approvals. Unusually rapid or repetitive approvals prompt targeted review rather than an automatic finding of misconduct. Depending on what that review shows, repeated failures to verify may require coaching, reassessment, restricted access, or incident review. Training cannot justify permissions or interface design that allows unsafe action.

The protected transaction record must show what was proposed, approved, attempted, and completed. It links that sequence to the patient, episode, source task, service version, coordinator, tool request, external response, and final outcome. It records action order and the receiving-system evidence used to verify results. A planned action, an accepted request, a confirmed action, verifiable message delivery, and any required patient or staff acknowledgement remain distinct. An uncertain outcome remains visible until a named person reconciles it; the agent's own success message cannot close the task. To identify work missing from the transaction record, reconciliation compares it with an independently maintained list of authorized source tasks. The monitoring record explains how both sources are checked for completeness and records any remaining gaps.

Automated monitoring reports changes to versions, settings, or permissions, transaction counts, failures, retries, prevented duplicates, waiting time in the queue, use of write-disable controls, and whether monitoring information is complete. It flags missing steps, partly completed actions, unacknowledged handoffs, and tasks approaching or exceeding their response deadlines. This reduces manual reporting as volume grows, but named owners still interpret findings and decide how to respond.

The monitoring plan requires a time-limited investigation when warnings arise, such as rising connector failures, a growing exception queue, increased rejection of proposals, or unusual approval patterns. If delay or workload approaches the point where timely follow-up or meaningful review may be compromised, the operational owner must restrict volume, add capacity, or return a defined set of actions to manual handling.

The hospital must immediately block writes and start incident response after confirmed approval bypass, cross-patient exposure, writes outside the allowlist, wrong-patient external actions, loss of independent write control, or inability to reconstruct transactions. Staff use the manual alternative and reconcile actions already sent or still in progress. If normal write disablement is unavailable or fails, the technical owner uses the tested independent means of blocking the connection or credentials. For each applicable measure, the record specifies the denominator, limit, reporting delay, authorized responder, and evidence needed for restart. Uncertain outcomes, missing required acknowledgements, and missing monitoring evidence receive the predefined response appropriate to their consequences and how quickly harm could occur. None counts as successful completion.

The committee reviews completed and failed actions, time to follow-up, and time spent waiting in the queue. It considers these together with manual corrections, complaints, subgroup findings, and coordinator workload. It also considers availability, supplier performance, and lifecycle cost when deciding whether the service still meets safety and value requirements. High transaction volume does not establish benefit or meaningful human

review. Because the service remains Tier 3, it continues through full review even if the hospital uses a streamlined pathway for eligible lower-risk services.

A3.7 Review changes and action permissions

An unchanged model identifier does not mean the agent's behaviour is unchanged. Its actions also depend on the software coordinating its steps, the action rules, and local controls. Suppliers or partners may change the model, coordinating software, or action policy. The hospital may change connectors, required data formats, tools, messages, permissions, memory arrangements, confirmation rules, or retry limits. Task mix, patients, coordinator practice, and the ability to handle queued work can also change without a software release. The approved record must capture these dependencies so staff can compare current transactions with authorized operation. An unexplained difference blocks further expansion until its source and effects are established and any required correction or renewed approval is complete. The owners also assess whether existing use must be restricted or paused; holding expansion does not contain a problem already affecting patients.

The hospital limits unapproved live changes by checking versions against the approval, restricting tools and credentials to the allowlist, and requiring supplier notice. Controls outside the language model enforce permitted actions. Monitoring checks for specified departures and records what it can and cannot detect. Updates are tested before entering the live environment, using transaction tests, deliberate attempts to defeat controls, and repeat workflow tests of previously validated behaviour. Material release changes require approval and production confirmation before deployment. Tested rollback and independent write disablement provide ways to contain a problem during investigation. If delay would create greater harm, an emergency safety change may use the separately authorized urgent route. The change must be recorded immediately and receive the required later review; mandatory duties still apply.

The party responsible for the changed component leads its correction. A supplier-controlled change is escalated under the agreed service requirements and may lead to deferral, rollback, restricted writes, correction, or exit. The hospital's integration team corrects connector, data-format, or required-field defects; affected writes stay disabled until records have been reconciled and repeat tests show that previously validated functions work. A change to shared action rules triggers the partner's evidence and rollback duties, while the hospital decides whether to release it locally. Changes in the patient population or task mix return to the clinical owner for evidence and approval review. Poorer coordinator confirmation or inadequate queue capacity may require restricted access, renewed demonstration of competence, more staff, or lower volume. The technical owner coordinates investigation while the cause remains uncertain. If restart or a material release change is required, passing tests is not enough: the required decision and production confirmation must also be complete before actions resume.

A3.8 Manage incidents, restart, and retirement

After a serious event, an authorized operator blocks the affected write actions and narrows or revokes credentials while preserving action and approval records. This control works independently of the agent and prevents further affected requests, including queued or scheduled ones. It does not establish what happened to requests already accepted by another system; those remain unresolved until their outcomes are checked. Staff return work to the manual queue, identify open and completed transactions, and reconcile appointments and messages with each receiving system. Each outstanding task has a named receiving person or service, with acknowledgement and escalation under the established response requirements. The team reviews affected patients and records, reverses or corrects actions where possible, and arranges required clinical follow-up and communication. Restoring an earlier software version does not undo a delivered message or completed booking; those effects need separate review and correction.

Investigators reconstruct the transaction from the signed instruction, patient link, and task information retained by the agent. They examine action rules, permissions, tool and connector responses, and coordinator review. They compare attempted and approved actions with actual appointments, messages, and records to find omissions, duplicates, partial or out-of-order actions, and actions using the wrong patient or task context. Queue conditions, supplier releases, partner changes, and recent history are tested against that evidence as possible causes. The component matrix assigns hospital, supplier, and partner duties. **Appendix H** records the chronology, affected actions, actual and credible harm, containment, evidence, notifications, causes, and corrective work, including lessons for subsequent practice. Responsible authorities complete any required safety, privacy, legal, regulatory, or patient notifications.

Before restart is considered, someone independent of the repair must verify objective evidence that each corrective action meets its defined pass criterion. Every action has a responsible owner. Applicable tests of transactions, resistance to deliberate attempts to defeat controls, and staff interaction with the service must pass again, as must tests of write disablement, manual fallback, and reconciliation. Monitoring and capacity to handle the exception queue must be restored. Tests cover both the original failure and related actions or setups it could affect. The hospital resolves affected records and actions and completes any required staff training or reassessment.

Once the responsible authority accepts the corrective evidence, the authorized body decides whether to permit restart and records the decision, allowed use, and restrictions. Production confirmation must be complete before the affected actions resume, and the restart record specifies closer monitoring and review. Restart does not authorize replay of old tasks or approvals. Outstanding work must first be reconciled, and every new consequential action needs current authorization for that individual case.

The hospital considers retirement if value is insufficient, safety or equity concerns remain unresolved, essential supplier or partner support is unavailable, or the service cannot stay within the number of configuration groups the hospital can support. Loss of recurring funding or a better replacement process may also justify closure. The decision-maker named in the hospital's authority matrix records the retirement decision with operational leadership, identifies the responsible owners, and specifies the replacement process.

The owners revoke identities and credentials and disable connectors and scheduled work. They reconcile open tasks and appointments, correct affected records, establish the outcome of requests already sent, and prevent queued or scheduled actions from executing later. Unfinished coordination goes to named staff who confirm acceptance, and the replacement process is restored. Patients and receiving services get the information needed to continue follow-up through that process. The owners complete contractual duties and verify required retention, return, or deletion of data and logs, including stored task information and memory. Retirement is recorded as complete only after transition, reconciliation, data, and contractual duties are completed and verified and no unsupported automation remains in operation. The registry keeps the closed entry as the service's history.

Findings from incidents and routine operation inform revisions to hazards, transaction tests, contracts, monitoring thresholds, competency scenarios, and configuration-group rules. A de-identified account may be shared through an authorized learning route where appropriate. The completed dossier records whether the second team may use the agent, the evidence for that decision, and the conditions that would require restriction, re-tiering, or retirement.

Appendix B Minimum governance model checklist

Hospitals with smaller AI portfolios can assign governance work to existing structures or use this checklist to establish a new governance program. Dedicated governance infrastructure is not required if these arrangements meet every required control and response time. The minimum model uses fewer standing resources but maintains the same safety standard. Staffing and support must increase when demand exceeds capacity.

Across the implementation checklists, an item is complete only when the evidence required for the current decision has been accepted and the verifier and date are recorded. The companion record uses Met, Open, or N/A. Each Open item has a status of Pending, Failed, Unknown, or Expired, together with the required action, owner, due date, and decision or activity it blocks. Evidence required at a later stage may remain Open—Pending while earlier activities that do not depend on it proceed. N/A requires a reviewer's approval of the reason; it cannot waive a mandatory control.

A full review examines all evidence that applies to the requested decision. A focused review examines a specified change or issue and checks that the other assumptions supporting approval still hold. These terms describe the extent of the review. For either type, evidence may be reused only after a documented comparison shows that it applies to the current use, context, version, and setup. The separately approved streamlined pathway is available only to qualifying Tier 1 or Tier 2 uses.

When several checklist items rely on the same report or other evidence, their records refer to one stored copy. Each completed item records what the evidence shows and the decision made; a note that someone reviewed it is not enough. If a requirement depends on a measure or service level, the linked evidence must identify the threshold and the response required when it is reached.

B1 Governance structure and authority

☐ The hospital has formally designated the AI Steering Committee or equivalent body responsible for AI governance. The designation and approved mandate are retained.

☐ The reporting record names the executive or executive leadership forum to which the AI Steering Committee reports. It states the reporting relationship and the committee chair's direct route for escalation.

☐ The Board's or equivalent governing body's endorsement of the hospital's AI risk appetite, risk tolerance, and Tier 0 screening criteria is retained.

☐ Each oversight body has an approved, retained mandate stating its purpose, scope, and reporting relationship.

☐ Each mandate states which decisions the body may make, its membership and quorum, how urgent reviews are requested, and who keeps its records. It also gives the effective date and next review date.

☐ The requirements for each type of review specify the expertise needed and the circumstances that require each specialist's involvement.

☐ Each review type has a named primary and alternate reviewer, with a contact route and maximum response time.

☐ When required expertise is unavailable, the affected decision is held or referred through the approved alternate-review route. The record states exactly which unrelated preparatory work may continue.

☐ The charter sets the maximum interval between routine committee meetings.

☐ The urgent-review procedure names who may request review and who may decide. It states the response deadline, required record, and executive or executive leadership forum for escalation.

☐ The authority matrix names who may assign risk tiers, approve a use and its conditions, and issue production confirmation—the recorded decision that the actual system meets the approval and all release conditions before live use. It also names who may authorize expansion, pause, restart, or retirement.

☐ For each decision, the authority matrix specifies delegation limits, escalation triggers, quorum requirements, and the time allowed for a response.

☐ The authority matrix states when a risk-tolerance exception, major incident, mission-critical tool, high-value contract, or unresolved risk remaining after controls must be referred to executive leadership. It also states that escalation cannot waive Tier 0 prohibitions or mandatory legal or specialist requirements.

☐ The capacity assessment compares current review capacity with the expertise and service levels required and records any gaps.

☐ The capacity plan assigns a remedy to each identified gap through local staffing, shared expertise, or regional support.

☐ Where regional review is used, the agreement records contacts, response times, decision boundaries, and rules for sharing records.

☐ The regional agreement has been checked to confirm that it identifies the decisions retained by the hospital and preserves its accountability for local use.

B2 Records and visibility of AI use

☐ The standard intake form identifies the proposed use, product and version, workflow, route followed by the data, enabled functions, need being addressed, owners, and requested decision. It also records the harm pathway, time to harm, detectability, reversibility, and scale of exposure, with the enforced limits and required human decisions used to define the use. It identifies the features or circumstances that trigger specialist review.

☐ The registry records one current lifecycle status for each use: Research only; Intake; Local assurance; Approved but not live; Early-adopter implementation; Routine operation; Restricted; Paused; or Retired. Any expansion decision is linked to the current status; expansion is not a separate status.

☐ Each use has a dossier. Its index links the description of the system and use being assessed (the baseline), supporting evidence, hazards, decisions, conditions, and records from each lifecycle stage. Each baseline is clearly marked as proposed or approved.

☐ The registry links the complete Tier 0 screen and the separate Tier 1–4 consequence assessment, including uncertainty and the authorized decision.

B2.1 Streamlined review

The hospital must approve a streamlined pathway before using it for eligible Tier 1 or Tier 2 uses. To qualify, a use must fall within the defined consequence limits, have a Tier 0 Proceed decision for the proposed activity, and have current supporting evidence and competent staff in the relevant roles. A use is ineligible while an unresolved finding requires another review route. Before the pathway opens, a baseline audit must meet the approved acceptance criterion. Later audits determine whether the pathway may continue or must be suspended while corrective work is completed.

☐ Before acceptance for streamlined review, the use is confirmed as Tier 1 or Tier 2, with a Tier 0 Proceed decision for the proposed activity. It must not materially affect care, access, rights, employment, an official record, or a safety-critical workflow.

☐ The use has been compared with the approved baseline.

☐ The evidence due for this decision is complete.

☐ Everyone performing the relevant roles has current competency evidence.

☐ The eligibility review checks for unresolved issues requiring specialist review, conditions that block the decision, serious incidents, and material supplier concerns. Any finding that makes the use ineligible for streamlined review results in referral to the required review route.

☐ An audit sampling plan is approved before the streamlined pathway opens. It identifies the eligible Tier 1–2 decisions and groups of uses, explains sample selection, defines a material error, names the independent reviewer, and sets the audit interval.

☐ Before streamlined review is activated, a retained baseline audit shows that comparable decisions completed through full review meet the approved acceptance criterion.

☐ Ongoing audits follow the approved interval and sampling plan. Findings are recorded separately by tier and use grouping.

☐ The sampling plan specifies which findings require suspension of streamlined review and who has authority to order it.

☐ Completion of corrective actions is verified before streamlined review is reactivated.

☐ Before streamlined review is reactivated, role-based competence is current.

☐ The follow-up audit meets the approved reactivation criterion before streamlined review resumes.

☐ The registry links the current tier, Tier 0 decision, whether data leave the hospital, privacy and security findings, supplier status, and monitoring plan.

B2.2 Governance staffing and software

☐ The capacity plan sets measurable thresholds for adding staff or workflow support. It states the current demand, how often demand and capacity are reviewed, and what action is required when each threshold is reached.

☐ The governance workflow has been checked to confirm that it tracks cases, decisions, versions, conditions, deadlines, monitoring alerts, supplier notices, incidents, and configuration groups. Each configuration group contains variants of the same defined use. Evidence must show that those variants can be evaluated, monitored, and supported together within approved operating limits.

☐ The governance-function inventory lists every function that automatically routes administrative work or offers a provisional recommendation and explains what each function does.

☐ Every final Tier 0, consequence-tier, residual-risk, and operational decision has a recorded authorized human decision-maker.

☐ The dossier links the value case, non-AI alternatives, data-flow map, approved operating limits, and plan for rollback or safely disabling the service.

B3 Lifecycle oversight

☐ The clinical/end-user, operational, and internal technical owners are named before any activity depends on their work.

☐ For a purchased or partner tool, the hospital's internal technical owner and the supplier or partner contact accountable for technical matters are named.

B3.1 Supplier accountability

☐ For each shared monitoring signal, the hospital–supplier responsibility matrix identifies who detects it, who reviews it, and who responds. It links the complete monitoring-measure and response records, including thresholds, deadlines, and evidence to retain.

☐ The procurement record distinguishes a time-limited vendor proof of concept (POC) from mandatory local assurance.

☐ The procurement record requires a separate vendor POC before purchase or commitment to a particular version and setup if material questions about suitability, performance, integration, or workflow fit cannot be answered by adequate existing evidence or an equivalent authorized evaluation.

☐ If a separate POC is omitted, the procurement owner and responsible governance authority record the evidence and reasons under local policy. Omitting a POC does not waive local assurance or authorize operational use.

☐ Before the initial limited rollout (early-adopter implementation), the record states the service need, specific use, and approved operating limits. It identifies the required user practices to be observed, first review date, authority to stop use, and separate decision needed for expansion.

☐ The supplier escalation plan covers missed deadlines, material updates, deterioration in service or performance, repeated failures, and missing evidence.

☐ For each trigger, the supplier escalation plan identifies whom to contact, when the response deadline starts to run, who may restrict or pause the service, and what decision is required.

B3.2 Local evaluation and incident readiness

☐ The version, setup, and data path proposed for live use have been tested in representative workflows in an authorized test environment.

☐ Any live activity has separate authorization.

☐ Acceptance criteria for performance, serious errors, relevant subgroups, usability, response capacity, and fallback are set before testing.

☐ Test results, limitations, deviations, and approval are retained and linked from the dossier.

☐ The screening record identifies the privacy and cybersecurity screening rule applied, reviewer, date, decision on further review, and reasons for that decision.

☐ Each privacy or cybersecurity condition is verified as met before the access or activity it governs is allowed. Every prerequisite for live use is resolved before production confirmation.

☐ For a use that could affect care, the patient-safety reviewer has provided a documented assessment. It covers credible ways harm could occur, time to harm, existing controls, and arrangements for detection and response.

☐ The safety reviewer's findings, residual-risk assessment, and conditions are retained.

☐ Each safety condition is verified as met before the decision or activity it governs. All conditions for live use must be met before production confirmation.

☐ Staff can use the incident-reporting route for the deployed service.

☐ Incident-response arrangements have been verified to allow staff to act within the time in which harm could occur.

☐ The incident-response record names who triages incidents, makes response decisions, and may contain or pause the service. It states the required deadlines and notifications.

☐ The latest incident-exercise results are retained. They show whether reporting, containment, and completion of the incident record met the approved criteria.

☐ Every failed incident-exercise criterion has an owner and due date for correction.

B3.3 Monitoring, change, and retirement

☐ Each monitoring measure has a recorded source, calculation, denominator, baseline, comparison window, minimum-data rule, and uncertainty rule.

☐ Each threshold has a recorded response, owner, deadline, and escalation route. Reaching a stop threshold requires immediate suspension of the affected function and activation of fallback.

☐ The restart procedure requires accepted evidence that the necessary corrections are complete and a recorded restart decision by an authorized person. Separate production confirmation is also required before the affected use resumes.

☐ Each Tier 3 or Tier 4 monitoring plan has been checked to confirm that it includes the applicable safety, performance, drift, subgroup, workflow, and incident measures.

☐ A material change has governance authorization before use in live operation, unless the separately authorized emergency-change procedure applies.

☐ Production confirmation for the changed version and setup is complete before live use, unless the separately authorized emergency-change procedure applies.

☐ The emergency-change procedure names who may authorize a temporary release. It specifies how harm will be contained, the deadline for reviewing the change afterward, criteria for rollback, and the record needed to justify continued use.

☐ Each deployed system has measurable retirement triggers, with a named decision owner and deadline for deciding.

☐ The retirement plan's arrangements for continuity, communications, record reconciliation, access removal, and return, retention, or deletion of data have been tested.

☐ Each completed retirement step has linked evidence of completion and a recorded verifier and verification date.

☐ Each unfinished retirement action has an owner and due date.

B4 Research and procurement responsibilities

☐ Research systems used in or connected to live workflows are registered so their use and operational dependencies are visible.

☐ Before a research system becomes an operational service, its handoff record states the limits of the research authorization and the proposed live version, setup, and use. It identifies evidence that can be reused and gaps that remain, operational owners, reviews that must be reopened, support and monitoring arrangements, and unresolved conditions.

☐ The research-to-operations handoff retains separate records of governance approval and production confirmation, the later decision confirming readiness for live use.

☐ The intake record states the proposed sourcing route and unresolved questions about components or vendor selection.

☐ The final build, buy, or partner classification and component responsibility matrix have been approved before commitment or approval of the proposed arrangement. That decision includes any relevant vendor-POC findings.

☐ The procurement review compares the exact product, version and setup, data path, and proposed local use with the approved procurement standard.

☐ The procurement review records evidence gaps, contract requirements, conflicts, exceptions, conditions, and the final decision.

☐ The contract assigns supplier duties for updates, monitoring, investigation of drift or degradation, and corrective action.

☐ The contract secures evidence and audit access, incident cooperation, and response and restoration deadlines.

☐ The contract sets limits on data use, prohibits unauthorized secondary use, and assigns retention and deletion duties.

☐ The contract specifies continuity arrangements, remedies, and exit support.

B5 Workforce, transparency, and learning

☐ Before receiving access, Tier 3 or Tier 4 users must demonstrate that they can follow the approved workflow, verify outputs, recognize prohibited uses, use the fallback, escalate concerns, and report incidents. These skills are assessed against defined pass criteria.

☐ Competency records show completion, failed attempts, the decision to grant or withhold access, and the events requiring reassessment.

☐ The literacy plan covers the Board or equivalent governing body, AI Steering Committee, Project Working Groups and system owners, frontline users, and patients or the public where relevant. For each audience, it states the required learning and assessment, how completion is recorded, and when education must be refreshed or competence reassessed.

☐ Governance members meet the relevant competency criteria before exercising authority.

☐ Tier 3–4 users meet the relevant competency criteria before receiving access.

☐ Each approved service has a record of the permitted product, version and setup, eligible users, permitted data, outputs or actions, setting, expiry date, and owner.

☐ Each service's approved workflow specifies the verification or approval step required before its output or action may be used.

☐ Prohibited uses are recorded separately from permitted uses, with the technical, access, and audit controls that enforce those prohibitions.

☐ Practice scenarios from the local portfolio cover normal use, serious failure, disagreement, fallback, and incident reporting.

☐ Each use has a record of whether disclosure is required and who has authority to make that decision.

☐ Where disclosure is required, the approved rule states whom to inform, what to tell them, when, and how the information will be made accessible. It also identifies practical human or non-AI alternatives offered where feasible or required, any limits on alternatives and how concerns will be addressed, and routes for questions or correction. A missing mandatory alternative prevents the affected use from proceeding.

☐ The disclosure record names who implemented the requirement, when it took effect, and the result of checking its implementation.

☐ A retained, dated exercise demonstrates that users can reach the incident-reporting route and on-call owner within the required time.

☐ The incident exercise confirms that registry and incident records agree.

☐ The incident exercise confirms that required evidence is preserved.

☐ The incident exercise demonstrates that the authorized person can suspend the affected function and activate fallback within the required time.

☐ Each missed incident-exercise target has a recorded corrective action, owner, and due date.

☐ The incident-sharing procedure names who authorizes sharing and who receives submissions. It specifies the minimum information to submit, where and when to submit it, and how feedback will be provided.

☐ For the most recent incident eligible for sharing, the decision to share or withhold it, reasons, and supporting reference are retained.

☐ The executive-reporting standard sets the maximum reporting interval. It requires reports to cover portfolio status, overdue reviews, threshold breaches, incidents, and material changes.

☐ Executive reports identify unresolved conditions, supplier failures, value findings, retirements, and the decisions or resources needed.

☐ The policy owner sets a maximum review interval and triggers for earlier review. Triggers include legal change, a serious incident, repeated threshold breaches, a new capability, or material portfolio growth.

☐ Each policy review has a record of the approved revision, effective date, responsible owners, and outstanding actions.

☐ The controlled list of locally adopted templates includes their source and approved deviations.

☐ A named owner reviews eligible de-identified lessons and incidents at the approved interval and submits suitable items through the authorized learning route.

☐ Each shared reference has a linked local decision to adopt or reject the resulting change, with the responsible owner, reasons, and date.

Appendix C Comprehensive governance model checklist

The comprehensive model applies when the portfolio's size, complexity, or recurring workload exceeds what the minimum arrangements can support. Every requirement of the minimum governance model still applies. The comprehensive model adds dedicated staff, infrastructure, and further assurance where workload and risk justify them. Hospitals using the minimum model must still provide baseline literacy and monitoring, along with any independent review required by risk, regardless of program size.

C1 Governance structure and capacity

☐ A standing AI Steering Committee operates under an approved mandate, with formal decision authority, quorum requirements, and scheduled meetings.

☐ An appointed governance lead coordinates intake, assigns reviews, and tracks the status of each AI use. The lead's authority, allocated time or capacity, reporting relationship, backup, and succession arrangements are recorded.

☐ The governance staffing record names the staff providing governance and analytical support and quantifies their allocated time or capacity. It states their service standards and backup coverage.

☐ Responsibility is assigned for registry quality, review preparation, monitoring reports, action tracking, and committee records.

☐ The requirements for each review state the disciplines needed, attendance expectations, designated alternates, and quorum rules.

☐ Attendance at each review is recorded.

☐ A decision is held or referred if a required reviewer or specialist finding is missing.

☐ Each lifecycle review has documented triggers, required inputs, and reviewers who must participate.

☐ The lifecycle review procedure states the available decisions, who may make them, and the maximum response time. It specifies the required record and conditions that prevent completion.

☐ The authority matrix specifies decision rights for each tier, delegation limits, and triggers requiring escalation.

☐ A sample of actual decisions has been audited against the authority matrix, and the findings are retained.

☐ Every departure from the authority matrix found by the audit has a recorded corrective action, owner, deadline, and evidence of completion.

☐ Measurable risk tolerances reflect the risk appetite endorsed by the Board or equivalent governing body.

☐ The annual review of those risk tolerances is current.

☐ The committee's delegated authority and escalation rules are documented. They state that escalation beyond risk tolerance cannot waive Tier 0 prohibitions or mandatory requirements.

☐ When using external benchmarking, the governance lead records the comparison source, measure definitions, comparison interval, and latest result. The record identifies material differences between the services compared.

☐ When a local result cannot be compared with the external benchmark, the record explains why and states the resulting decision, action owner, and due date.

C2 Records and supporting systems

☐ The registry has been checked to confirm that it supports the shared governance record fields, search, version history, and lifecycle-status tracking. It also covers external data flows, supplier changes, equity monitoring, and retirement.

C2.1 Governance workflow platform

Workload thresholds determine when dedicated staff or a governance workflow platform must be added. When demand reaches the approved platform threshold, the hospital must have a working system that routes reviews, tracks conditions and limits on configuration groups, receives alerts, and keeps an exportable decision record. Each configuration group contains variants of the same defined use; evidence must show that these can be evaluated, monitored, and supported together within approved operating limits. The platform may automate administration and offer provisional recommendations, but authorized people remain responsible for substantive decisions.

☐ Measurable workload thresholds for adding dedicated staff or a governance workflow platform are approved.

☐ When demand reaches the approved platform threshold, the platform is verified as in place and working. It routes reviews, tracks conditions and limits on configuration groups, receives alerts, and retains an exportable decision record.

☐ Platform automation is limited to administrative routing and provisional recommendations. Authorized people make all substantive decisions.

C2.2 Record maintenance and reporting

☐ The template-control procedure lists the current required templates and their effective dates for each applicable lifecycle decision.

☐ A sample audit records use of current and outdated templates and any deviations.

☐ Each corrective action from the template audit has a recorded owner and status showing whether it is still open or complete.

☐ The dossier retains evidence and decisions for all six recurring tasks: assign accountability and risk limits; define the use and expected value; review data and sourcing; assess readiness for use; review operation and respond; and share learning and manage retirement. These tasks are applied wherever needed across the seven decision stages.

☐ The portfolio-dashboard standard names who prepares and reviews the dashboard, how often it must be reviewed, and the measures it must contain. Reports must identify overdue items and record the decisions made and actions taken.

☐ The executive-reporting record retains the approved schedule, latest issue date, recipients, and any overdue escalation.

☐ The change-control log identifies model, prompt, data, workflow, and vendor changes and links them to their review and decision records.

☐ The library of standard contract clauses and the vendor scorecard each have an assigned owner.

☐ The contract-clause library and vendor scorecard have controlled versions, a maximum review interval, and triggers for earlier review.

☐ Each clause-library or scorecard revision retains the evidence sources, approval, effective date, and superseded version.

☐ Standard operating procedures (SOPs), training materials, and approved disclosures for users or patients are centrally accessible and version-controlled.

C3 Evaluation, monitoring, and research

C3.1 Standards for evaluation and oversight

☐ The local-validation standard specifies when validation is required according to the risk tier and credible pathways to harm.

☐ The local-validation standard specifies the required study design, the reference used to determine correct results (reference standard), and what the study must cover. It sets criteria for performance, serious errors, relevant subgroups, and uncertainty.

☐ Each departure from the local-validation standard records the approver, reason, and evidence addressing the resulting gap.

☐ The evaluation standard defines when subgroup, equity, and human-factors reviews are required. For each applicable review, it specifies the method, reviewer, and decision rule.

☐ Default thresholds are maintained for applicable data drift, performance, calibration, subgroup harm, alert burden, unsafe generated output, and action failures. Any adjustment for a particular use is justified by the credible ways harm could occur in that use.

☐ For each Tier 3 or Tier 4 use, the dossier records the assurance approach and explains how possible consequences, time to harm, workflow, evidence gaps, and practical feasibility influenced the choice.

☐ When an expected assurance method will not be used, the approver's recorded reason and description of the alternative evidence required are retained.

☐ Research systems are tracked in the registry.

☐ Research systems complete the formal research-to-operations pathway before becoming operational services.

C3.2 Routine monitoring and reassessment

☐ Each routine monitoring measure uses the shared monitoring record, including the version of the calculation method, denominator, comparison window, and rule for the minimum data required.

☐ The shared response record states the warning, action, stop, and restart criteria, required responses, responsible people, deadlines, and escalation routes. Reaching a stop threshold requires immediate suspension of the affected function and activation of fallback.

☐ Each Tier 3 or Tier 4 use has a maximum reassessment interval and defined events that require earlier review.

☐ Reassessment compares the current service with its approved baseline and reviews hazards, monitoring, incidents, changes, user practice, subgroup findings, value, supplier performance, and retirement readiness.

☐ Each change prompted by monitoring has links to the records or materials it affects and a record of its owner, deadline, approval, effective date, and evidence of completion.

☐ When an authorized reviewer decides that a monitoring finding needs no action, the evidence, reasons, and decision authority are retained. This decision cannot override a mandatory stop response. Any emergency continuity decision needs separate authorization and documentation.

☐ The retirement procedure defines measurable triggers, the authority responsible for the decision, and the deadline for deciding.

☐ The retirement plan explains how work will continue, who will be informed, how access and integrations will be removed, and how records will be reconciled. It also covers how data and remaining contract obligations will be handled when the service ends.

☐ Completion of the retirement work has been verified and the result recorded.

☐ The registry has been updated to reflect retirement.

☐ Each retirement lesson or unresolved action has an owner.

C4 Operations, contracting, and resilience

☐ The sourcing assessment compares build, buy, and partner options using approved criteria and weights for fit, evidence, control, capacity, lifecycle cost, continuity, and exit.

☐ For each material component, the sourcing record documents the selected route, any exceptions, and the reasons for the decision.

☐ The sourcing record identifies the scorecard and clause-library versions used, deviations, and approved exceptions.

☐ Each unresolved contractual gap has an owner and deadline.

☐ A sample of procurements has been audited to confirm use of the current scorecard and clause-library versions.

☐ Supplier scorecards assess the adequacy of update disclosures, monitoring and investigations, audit access, and cooperation during incidents. They also assess corrective action, service levels, and whether the hospital can leave the arrangement in practice.

☐ An assigned contract owner audits compliance with supplier obligations at the approved interval and after the specified events.

☐ Each supplier-compliance audit compares supplier evidence with the contract terms and hospital monitoring.

☐ Every breach of a supplier obligation has a recorded remedy, responsible owner, deadline, escalation route, and evidence that the required action is complete.

☐ A vendor change, late disclosure, unresolved degradation, or repeated failure has a formal change record.

☐ Each supplier change has a record of the affected version and setup, interim restriction, reviews to reopen, deadline, and who may decide. The record also states the rollback or exit action and retains evidence that the required work is complete.

☐ The contract defines which changes are subject to its change-control terms, the required notice periods, and when further review is triggered. It sets uptime and recovery targets, incident deadlines by severity, and how often continuity must be tested. It preserves remedies and the hospital's rights to defer, restrict, or exit.

☐ The asset inventory covers data, logs, prompts, templates, outputs, derived data, and model or configuration artifacts.

☐ Each asset has a linked governing contract clause in the dossier and a record of the rights held by each named hospital or supplier party.

☐ Tests at the approved interval establish whether the agreed rights of access, export, deletion, restriction, and audit can be exercised.

☐ Each breach of those rights has a recorded remedy and evidence that the required corrective action is complete.

☐ The technical owner compares approved data flows with observed flows at the approved interval and after defined architecture changes.

☐ The data-flow audit retains its evidence and findings. Its record identifies the affected data, how the problem was contained, how the data were handled, how discrepancies were resolved, and evidence that the required actions are complete.

☐ Supplier changes to models, prompts, guardrails, data pipelines, and infrastructure have undergone impact screening.

☐ Each material supplier change has completed the required change-control process before release.

☐ One record tracks required specialist reviews throughout the lifecycle. For each review, it shows whether it applies, its owner and status, findings and conditions, and evidence that the conditions have been met.

☐ Any decision or activity affected by an unresolved mandatory specialist finding is held. Every mandatory finding required for live use is resolved before production confirmation.

C5 People, transparency, and improvement

☐ The role-based competency matrix states what each audience must know or demonstrate, how and when it will be assessed, and the pass criteria. It names who keeps the records and states when education must be refreshed or competence reassessed. It also explains what happens when someone does not meet the criteria.

☐ The transparency record states which notice, complaint, correction, explanation, and human-review routes apply to the use.

☐ For each applicable route, the transparency record identifies the decision-maker, response deadline, practical human or non-AI alternatives offered where feasible or required, and latest test result. It records any limits on alternatives and how concerns will be addressed. A missing mandatory alternative prevents the affected use from proceeding.

☐ For each Tier 3, Tier 4, or sensitive use, the engagement record states the decision on patient, staff, or community engagement and its rationale. It identifies who participates and the support provided to enable participation.

☐ The engagement record documents the issues participants raised, changes made or declined, and the reviewer accountable for considering their input.

☐ An annual audit examines a sample of uses selected according to risk. It traces each use from intake through its current monitoring, change, incident, and retirement-readiness records.

☐ Annual-audit gaps have assigned owners and deadlines, with retained evidence of closure.

☐ The incident-learning record identifies events eligible for sharing, corrections made, de-identification, decisions on sharing, and references to the material shared.

☐ Each recommendation received has a record of its adoption or rejection, reasons, responsible owner, and date. Evidence that any resulting action is complete is retained.

☐ The annual maturity assessment is complete. The assessment tool and version, supporting evidence, reviewers, results, and changes in results over time are retained.

☐ The improvement plan assigns funded actions, owners, and deadlines to priority gaps identified by the maturity assessment.

☐ Executive acceptance of the improvement plan and the next assessment date are recorded.

☐ The resource collection has a named owner and review interval.

☐ Relevant updates to the resource collection are logged.

☐ Each relevant update has been assessed, with its local adoption or the reason for a deviation recorded.

☐ Each adopted update has a record of its approval, effective date, and superseded local material.

☐ Each maturity or survey finding is linked to a local priority or has a recorded reason why no change is needed.

☐ Each local response has a record of the resources, owner, target date, and result when the work is completed.

☐ Each de-identified finding shared through the authorized learning route has a recorded sharing authority, date, and reference.

Appendix D Evidence requirements by risk tier

This checklist covers one defined AI use. The Tier 0 screen is completed first. After a Proceed decision, the assessment covers the risk-tiering factors and the core requirements for Tiers 1–4, together with the requirements for the assigned tier and all lower consequence tiers. Evidence is due at the decision it supports. The system's capabilities or mandatory specialist findings may require more assurance than the tier list specifies. An item is marked N/A only when it does not apply and a reviewer accepts the reason.

D1 Tier 0 Initial screen

Tier 0 assesses whether the next proposed activity may begin; it does not assign a consequence tier. Screening covers law and policy, authorized data and system access, privacy, cybersecurity, and required approvals. It addresses foreseeable harm, rights, equity, accessibility, cultural safety, and applicable Indigenous-led governance, alongside human review, auditability, fallback, supplier cooperation, oversight evidence, change notices, incident support, and exit rights. The record identifies the baseline used to assess benefit and the minimum worthwhile improvement, or names who will establish them before the relevant evaluation. A named person approves safety, performance, and equity criteria before results are reviewed and before the relevant activity starts. Unapproved use is stopped or restricted before review continues, with any separate, time-limited permission to continue linked. Evidence due later may remain pending until the decision that requires it.

☐ The requested activity and its exact limits are recorded.

☐ The hospital owners responsible for the activity and the person authorized to stop it are named.

☐ Screening identifies any prohibited use or unacceptable feature that cannot be corrected.

☐ The detailed Tier 0 screening assessment is complete, with linked evidence and any missing essential information identified.

☐ The screening record states whether the permissions and safeguards required before this activity starts are in place.

☐ Each specialist review has a recorded completion deadline and the decision requiring its finding is identified.

☐ Each unresolved requirement has a recorded action, owner, due date, and activity that must wait.

☐ The Tier 0 decision is recorded with its reason, decision-maker, date, and authority for reconsideration. It states the permitted activity, limits, owner, expiry, and requirements for closing each open condition. Any action to stop or restrict current use is recorded separately.

The recorded decision is **Reject** for a prohibited use or one with an unacceptable feature that cannot be corrected, and **Hold** if essential information or a prerequisite for the requested activity remains unresolved. A **Proceed** decision permits only the specified review or evaluation under its recorded conditions; it does not authorize live use.

D2 Risk-tiering factors

D2.1 Consequence and exposure

☐ The tiering record traces the most serious credible consequence of the proposed use: how it could occur, how quickly harm could follow, who could be exposed and at what scale, whether it could be detected, and whether it could be reversed. Only limits and human decision points that are implemented, enforced, and verified count toward the assessment.

☐ The risk remaining after controls is assessed and recorded separately from the consequence tier.

☐ The consequence assessment links its supporting evidence and explains uncertainty. Proposed or untested safeguards do not justify a lower tier.

A consequence is credible when it could plausibly arise from intended use or reasonably foreseeable misuse. That pathway needs evidence or an explicit account of what remains uncertain; imagining an outcome alone is not enough. If material uncertainty prevents reviewers from ruling out a higher consequence tier, that tier is assigned provisionally and the evidence needed to resolve the question is identified.

☐ The tiering record identifies effects on care, access, rights, employment, resources, communication, records, privacy, security, equity, and workflow. For each affected area, it records the decision or action and who makes or receives it.

☐ The tiering record describes any effect on an official record or patient communication.

☐ The credible consequences of an error, delay, or inappropriate reliance on the system are recorded.

☐ The record lists the decisions or actions the system may influence or execute and identifies each point where human approval is required.

☐ The review record specifies the information, time, competence, and authority needed for each required human review, and states the practical alternative available.

☐ The record describes how users can override the system and links workflow-test evidence showing that required approval and override both work.

☐ The proposed deployment scale is recorded, including the number of patients, staff, sites, and transactions that may be affected.

☐ The equity assessment identifies locally relevant groups and explains how unequal error, access, burden, or benefit could arise for them.

☐ Available subgroup evidence is retained, with uncertainty and material evidence gaps recorded.

☐ Each required action to reduce risk has a named accountable owner and a defined finding that would require restriction or further review.

D2.2 Change and dependency

☐ The record states how often each material component may change and who may initiate and approve a change.

☐ The change profile specifies the notice, evidence, and local testing required for each component change. It states when the system must be rolled back or its tier reassessed.

☐ Any component that can change without the hospital's approval is identified.

☐ Every integration that reads or writes data and every operational dependency has a continuity record explaining how work will continue if it fails or becomes unavailable.

☐ The continuity record names who may disable and restore the service. It states the permitted outage limits, recovery time, and approved version and setup to restore (rollback target), or the tested plan for safe disabling and continuity if rollback is unavailable.

☐ The latest continuity-test result and any corrective action are retained. If a criterion is failed, the dependent activity remains on hold until the required evidence has been accepted.

D2.3 Patient-facing uses and external data

☐ The tiering record states whether the system is patient-facing or generates patient communications.

☐ The tiering record states whether personal health information or other sensitive data leave the hospital.

D3 Core requirements for Tiers 1–4

D3.1 Owners and approved use

☐ The clinical/end-user, operational, and internal technical owners and, where required, the executive sponsor are named. Each role's authority, accepted duties, backup, acceptance date, and contact route are recorded.

☐ Every material component is classified as Build, Buy, or Partner and has a named hospital owner.

☐ The component responsibility matrix assigns responsibility for design, changes, operation, maintenance, access to evidence, and response to problems.

☐ The component responsibility matrix records service expectations and the arrangements and dependencies for funding, backup, continuity, and exit.

☐ The defined use has a current AI registry entry.

☐ The proposed version is recorded during review and the approved version after the decision. Their status is clearly labelled so that the proposed version cannot be mistaken for an authorized release.

☐ The intended use, approved users, and setting are recorded, with out-of-scope uses listed separately.

☐ The initial privacy and cybersecurity assessment confirms the authority to use and access the data, identifies external parties, and defines when further review is required.

☐ Required privacy and cybersecurity assessments have a recorded screening rule, reviewer, date, decision on further review, and reasons. Their findings and conditions are linked to the decisions or activities they govern; each condition has an owner, due date, authorized verifier, and evidence accepted to close it when available.

☐ The initial patient-safety assessment records how harm could credibly occur, how quickly it could occur, which groups could be exposed, and what detection and response are needed.

☐ Every required safety review is assigned, and its findings and conditions are tracked against the decision they support.

D3.2 Monitoring and transparency

☐ Every material safety, performance, data, subgroup, workflow, value, or reliability risk has an assigned monitoring measure.

☐ Each monitoring measure has a recorded definition, source, denominator, calculation version, baseline, and comparison period. The record states the minimum data required, how uncertainty will be handled, any required reporting by site or subgroup, and the response to missing or late data.

☐ Each measure has a named primary and backup reviewer, review interval, and report recipient. The last and next review dates and the escalation required if review is overdue are recorded.

☐ Each measure's warning, action, stop, and restart criteria are linked to a response record stating the required action, authorized primary and backup responders, deadline, and escalation route. The record explains how uncertainty, insufficient data, or reporting delay affect the response, and specifies fallback, communications, any required review of affected records, and evidence needed to continue under restriction or restart. Meeting a stop criterion requires immediate suspension of the affected function, activation of fallback, and incident escalation. Restart requires accepted evidence, authorization for the scope of use, and production confirmation.

☐ The transparency record states whether people must be informed about the AI use and explains why. When disclosure is required, it specifies how the information will be provided, the alternatives available, and routes for questions, correction, and human review.

D4 Tier 1 Low risk

☐ A short factsheet explains the intended use, its boundaries, and limitations.

☐ When patient data or an external service is involved, privacy and security reviews cover the actual data path and architecture. The dossier links the findings, required controls, risk remaining after controls, and evidence that conditions for the proposed activity have been met.

☐ Named intended users have completed scenarios for normal use, errors, escalation, and fallback. Their performance has been assessed against criteria set before testing.

☐ The scenario record documents task completion, errors, delays, accessibility barriers, unresolved issues, and corrective action, with owner sign-off.

☐ The Tier 1 monitoring plan identifies the indicators that matter for ongoing oversight and the changes that require review, using the core measure and response fields above. It records how to report incidents and who may restrict or retire the use.

D5 Tier 2 Moderate risk

☐ The sourcing record explains the selected build, buy, or partner strategy.

☐ The evidence has been assessed for applicability to the exact release, use, population, setting, and workflow, and the reviewers' finding is retained.

☐ The record identifies gaps in performance evidence, limitations, and support and update arrangements. It records the reviewers' decision to accept the evidence, accept it with conditions, require further work, or reject it.

☐ The dossier describes the proposed workflow and its accountable owners.

☐ The dossier records whether local validation is required, who made that decision, and why. Where validation is required, its plan, acceptance criteria, results, limitations, and resulting review decision are linked.

☐ The dossier records whether privacy, cybersecurity, legal, and procurement reviews are triggered.

☐ Every required review has a retained record of the reviewer's identity, findings, mandatory conditions, and sign-off.

☐ Each unresolved issue has a named owner, due date, and decision it blocks. Once the authorized reviewer accepts evidence that the issue is resolved, that evidence is linked to the record.

D6 Tier 3 High risk

☐ Acceptance criteria for overall results, serious errors, and subgroups are set before evaluation begins.

☐ The proposed production release, applicable decision thresholds, and data path have been tested in the intended population and workflow.

☐ The local-validation report retains the results, uncertainty, deviations, limitations, and decision resulting from the review.

☐ Representative users have completed scenarios covering normal and difficult work, disagreement, errors, escalation, and fallback. Their performance has been assessed against criteria set in advance.

☐ The human-factors report records errors, delays, over-reliance, workarounds, and accessibility barriers. It includes corrective actions and the decision resulting from the review.

☐ The Tier 3 equity review identifies relevant groups and measures and explains uncertainty and missing evidence.

☐ The equity review applies the pre-set decision rules and records the mitigation, responsible owner, and review decision.

☐ The Tier 3 evidence plan links the validation and human-factors work. It specifies the comparator, representative data and workflows, checks of the local data and processing pipeline, limitations, and events that require reassessment.

☐ Before expansion is approved, the review examines evidence from the initial limited rollout (early-adopter implementation), including how well it represents the proposed wider use, adherence to required user practices, and capacity for support and monitoring. Each proposed change in site, population, language, volume, feature, integration, action, or autonomy is assessed separately.

☐ The Tier 3 monitoring plan contains complete measure and response records, including each measure's definition, source, calculation, denominator, baseline, comparison period, minimum data, uncertainty rules, required site or subgroup reporting, reviewers, review dates and intervals, and report recipients. It specifies thresholds, authorized responders, deadlines, escalation, and responses to missing data or overdue review. The plan states how work will continue during an outage and what evidence and authorization are required before restart, including production confirmation.

☐ The incident arrangements document the reporting route, deadlines by severity, communications, notifications, and authority to contain or pause the service.

☐ The Tier 3 continuity record identifies the approved rollback target or plan for safely disabling the service, latest test result, restart criteria, decision authority, and required evidence.

☐ A new production confirmation is obtained before restart.

☐ The Tier 3 change-control record identifies material-change triggers, pre-release approval authority, and the emergency-change route.

☐ An applicability matrix records which specialist sign-offs, competency checks, and disclosure actions are required.

☐ Each applicable item has a recorded authority, responsible person, date, result, any condition still to be met, and evidence that completed conditions have been met.

☐ Production confirmation is withheld while any mandatory finding required to confirm readiness for live use remains unresolved.

☐ For a use that materially influences care, the dossier records whether prospective evaluation with outputs prevented from influencing care (silent evaluation) is feasible and necessary. It names the decision-maker and explains the decision.

☐ Where prospective silent evaluation is used, the record documents how outputs are prevented from influencing care. The sample, duration, criteria, findings, and resulting decision are retained.

☐ If prospective silent evaluation is not used, the dossier identifies the alternative assurance method and evidence.

D7 Tier 4 Very high risk

☐ The Tier 4 assurance plan addresses the credible ways harm could occur and has governance approval before evaluation begins.

☐ All required Tier 4 assurance is complete before production confirmation.

☐ The assurance methods and representative scenarios, acceptance thresholds for serious errors and subgroups, and criteria for how people and the AI system work together are set before evaluation begins.

☐ An independent review records the results, limitations, and decision to pass or hold. If a method cannot be used, the reasons are retained and the substitute used is identified.

☐ Any proposed departure from an approved assurance method has a record of the unmet criterion, result, reason, and proposed substitute evidence or equivalent control. The departure must not waive a Tier 0 prohibition or mandatory safety, legal, or specialist requirement.

☐ The authorized exception decision is retained, including evidence of equivalent protection, conditions, effective date, expiry, and next review. Live use remains blocked when an applicable critical criterion has failed.

☐ The Tier 4 plan sets the longest permitted intervals between monitoring checks and governance reports. It also specifies any required real-time alerts or alerts triggered by an event.

☐ The Tier 4 plan names alert recipients and states how overdue work is escalated. It identifies who may investigate, restrict, or pause the service and their response deadlines.

☐ Before production confirmation, an exercise of fallback and either rollback or the approved plan for safely disabling the service has demonstrated that the approved continuity and recovery criteria are met.

☐ Before production confirmation, the incident-response exercise meets its approved criteria.

☐ Before production confirmation, the retirement exercise meets its approved criteria.

☐ The exercise record contains the results, defects found, owners responsible for addressing them, retest dates, and the decision to proceed or hold.

Appendix E Lifecycle implementation checklist

This checklist follows one defined AI use from the initial proposal to retirement. A checked item confirms that the action is complete or the stated condition has been verified, with supporting evidence linked in the dossier. The evidence record identifies the verifier and verification date. Every prerequisite must be met before the decision or activity that depends on it. Evidence needed for later evaluation or implementation need not be available at intake, but must be complete before the decision or activity that requires it. Hospitals using the minimum or comprehensive model may allocate the work differently; the decisions and controls required for each use remain the same.

The hospital-wide duties apply to every sourcing arrangement and are reviewed first. The build, buy, or partner checks then apply to each component according to how it is developed or supplied. Each duty must be completed before the decision or activity that depends on it. Different components may use different routes; no component needs to use all three. Checks involving supplier contracts or external contacts apply only where those relationships exist. Work performed entirely within the hospital does not require a supplier contract.

E1 Proposal and expected value

E1.1 Proposal and ownership

☐ The proposal describes the need, the proposed use of the product, and the decision requested. It records the exact product or version, intended users and setting, and prohibited uses.

☐ The proposal shows where the system enters the workflow, which systems it connects to, and how data will move. It identifies how the system will be developed or supplied, how independently it may act, the responsible owners, proposed timing, and known constraints.

☐ Every answer in the risk-tier record links to its supporting evidence.

☐ Reviewers have considered Tiers 4 to 1 and assigned the tier that reflects the most serious credible consequence of the proposed use. Any operating limit or required human decision used to define that use is already implemented, enforced, and verified. Planned or untested safeguards have not been used to justify a lower tier.

☐ If the evidence cannot exclude a higher tier, that tier is assigned provisionally. The record explains what remains uncertain and what evidence is needed to resolve it.

☐ The Tier 0 record states Reject, Hold, or Proceed and links evidence for every criterion. It identifies failed or unresolved conditions, the next activity requested, and any interim work permitted. It also states what must change before the decision can be reconsidered.

☐ The clinical/end-user, operational, and internal technical owners have accepted their duties and authority, with agreed backup arrangements. No activity proceeds if a role it depends on is vacant or the responsible person has not accepted the duty.

E1.2 Initial specialist screening

☐ Privacy and cybersecurity screening has determined whether the proposed data use and external connections need further review. The screening record identifies the data and connections assessed, the rule and version applied, the reviewer, date, and reasons for the decision.

☐ Each required privacy or security review, access-control condition, or logging condition has an owner and due date. The record states which activity must wait for completion and links the evidence accepted to confirm that the requirement has been met.

☐ Patient-safety screening covers effects on care, access, clinical communication, official records, and time-critical work.

☐ The safety assessment records the credible ways harm could occur and how quickly it could develop.

☐ Every required review has a named reviewer and due date.

☐ The governance lead records the Tier 0 decision, decision-maker, authority, date, and route for further review. The record specifies any redesign or additional information required before reconsideration.

☐ Where the tier remains uncertain, the triage decision states what is unresolved and exactly which activities may continue while it is investigated.

☐ Each condition arising from triage has a named owner, due date, and specified evidence. The decision on whether the condition has been met is recorded.

E1.3 Need, value, and hazards

☐ Each objective has a value record stating its measure, data source, denominator (the total quantity used to calculate a rate), baseline, target, and target date. The record includes relevant subgroup analyses, the owner, review interval, and action required if the target is missed.

☐ The baseline record describes the current service's process, users, volume, and timing; safety and quality results; workload; and access. It includes the relevant subgroup results.

☐ The baseline record identifies the data source and observation period, known points of failure, and the non-AI process or service used for comparison when judging the proposal's value.

☐ For each hazard, the hazard assessment explains what could fail, how it could cause harm, the consequences, who would be exposed, and how quickly harm could occur. It records the controls intended to prevent the hazard and the signals that would reveal it.

☐ Each hazard has detection thresholds linked to the required response, responsible owner, and approving authority. Meeting a stop criterion requires immediate suspension of the affected function, activation of fallback, and incident escalation.

☐ The hazard assessment records the risk remaining after controls.

☐ A retained comparison assesses the current process, feasible non-AI improvements, and AI options against the approved value criteria.

☐ The comparison records the option chosen, the reasons for choosing it, and why each material alternative was rejected.

☐ The authorization specifies the product and release, users, population and setting, data and workflow, and scale. It states how the output may be used or relied on and the duration of use.

☐ The authorization lists prohibited users, data, settings, outputs, actions, and downstream uses separately from what is authorized.

E2 Development and procurement

E2.1 Hospital responsibilities

☐ The sourcing decision compares Build, Buy, and Partner against the approved criteria: fit with the intended use, strength of evidence, control, capacity, total lifecycle cost, continuity, and exit. The comparison supports the route chosen for each component; the service need not use all three routes.

☐ For every material component, the sourcing record gives the comparison result, reviewers, reasons for the chosen sourcing route, and any approved exception.

☐ Every decision the hospital must retain has a named owner and backup before the hospital commits to an arrangement.

☐ For each lifecycle duty, the responsibility record identifies who does the work and who has decision authority. It states the service to be provided, evidence source, funding, and escalation route.

☐ The responsibility record states what happens if a lifecycle duty has no owner or is not performed.

☐ A named internal technical owner is accountable for the technical evidence submitted for governance review. This includes assessing the architecture and whether the proposed validation can be carried out.

☐ The internal technical owner is responsible for interpreting monitoring results and coordinating the response to technical incidents.

☐ The authority matrix states which technical findings, evidence-review findings, incidents, supplier issues, or remaining risks must be referred to the AI Steering Committee, and by when.

☐ The data-flow map has an identifiable version and shows every data source, element, transfer, recipient, connection point (endpoint), and storage or log location. Its record identifies hosting regions, suppliers, and subprocessors, whether personal health information leaves the hospital, retention and deletion arrangements, any permitted further use, and the authority for access.

☐ A dated comparison of actual flows with the data-flow map records any differences.

☐ The rights register covers source data, derived data, model files and related artifacts, code, logs, prompts, templates, and outputs. For each category, it identifies the rights holder and hospital custodian and states who may access, change, export, or delete it. It also records retention requirements, audit evidence, and the procedure for disputes or exit.

☐ The architecture and data-path review records why the review was needed, who performed it, and the reasons for the decision. For each required assessment or control, the record links the findings, conditions, deadlines, and evidence accepted to confirm completion.

☐ Reviewers assess whether the available evidence applies to the exact release, population, setting, workflow, and outputs proposed for use. They record their assessment.

☐ The sourcing review identifies gaps in evidence of performance, documented limitations, arrangements for updates, and assurance needed at the hospital. It records whether to accept the proposal, impose conditions, seek more evidence, or reject it.

☐ Every performance claim records what was measured, the acceptance threshold, and the exact release and configuration tested.

☐ Each claim links to its result, including the dataset and period, population and required subgroups, comparator, and uncertainty.

☐ The evidence record documents known limitations and circumstances that could reduce performance, with their evidence sources. It names the reviewer and date and records the decision to accept, set conditions, or reject.

☐ The responsibility matrix identifies each component's sourcing route, the party that controls it, the hospital owner, and the external contact where one is involved.

☐ For each component, the responsibility matrix assigns the work of design, change, operation, maintenance, evidence provision, monitoring, failure investigation, correction, and rollback or safe disabling. The allocation reflects which party controls the component.

☐ The responsibility matrix records the response times, funding, backup and continuity arrangements, and the dependencies that must be addressed to exit each component arrangement.

☐ Each signal of a material decline in service or performance, or a safety concern, has a measure record identifying what is measured, its source, denominator, calculation version, baseline, comparison period, and current result. The record states minimum data requirements, uncertainty and missing-data rules, required site or subgroup reporting, and review and reporting arrangements. It links the thresholds and required responses and names the person or system that detects the signal and the hospital reviewer who assesses it independently.

☐ Each material signal has a named first responder, investigation lead, and corrective-action owner. The responsibility matrix sets the response deadline and escalation route.

☐ Where several parties must respond, the responsibility matrix names the lead and the person or body that resolves disputes. It specifies how work and required information pass between parties and the deadline for sharing evidence. It also identifies who at the hospital may pause the service.

☐ The hospital has verified that it can inspect or challenge each safety-critical signal supplied by a vendor independently. It uses its own control or an approved alternative control that provides the required protection.

☐ The same responsibility matrix records shared duties. Each shared duty has a named lead, hospital decision authority, and external performer, with a specified handoff, evidence requirement, and deadline.

☐ For every supplier or performance problem requiring escalation, the procedure defines when the response clock starts. It names primary and backup contacts and sets deadlines for acknowledgement, action, and restoration.

☐ For each escalation trigger, the procedure identifies who may restrict or pause the service, the available contract remedies, and where an overdue response is referred.

☐ For each proposed update, the procedure requires identification of the affected component, exact change, owner, and person authorized to approve the change. It specifies the impact screening, evidence, and acceptance criteria required.

☐ The update procedure specifies when notice must be given and when the change may be released. It also sets requirements for monitoring after the change, rollback or safe disabling, and handling an emergency change, including the triggers for those actions.

☐ The incident agreement names primary and backup contacts, sets notice and response times for each severity level, and assigns responsibility for preserving evidence.

☐ The incident agreement assigns responsibility for investigation, communication, containment of harm, pausing the service, corrective action, and regulatory support. It also assigns responsibility for completing and recording incident closure.

☐ The evidence-access requirements specify which logs, documents, and evaluation materials must be available, how long they must be retained, who may access them, and the deadlines for delivering them.

☐ A dated access exercise has confirmed that named hospital investigators can obtain required evidence within contract deadlines. The exercise covers logs, records of system versions and configurations, evaluation and incident evidence, and corrective-action records. It also covers records of where data and system components came from and how they changed.

☐ If required evidence is unavailable, the hospital records its decision about the remaining risk and any equivalent control it accepts. The activity that needs that evidence remains on hold if its absence prevents an assessment or a mandatory finding.

E2.2 Supplier responsibilities and exit

☐ The supplier's current subprocessor register lists organizations it engages to process data on its behalf as part of the hospital service. Other external suppliers and data recipients are recorded separately. Each organization's entry records the service provided, data accessed, purpose, location, retention, and any further transfers.

☐ The contract requires advance notice of subprocessor changes and preserves the hospital's right to object or exit.

☐ The contract allocates liability and indemnity for patient or worker harm, privacy or cybersecurity events, intellectual-property claims, service failure, and regulatory non-compliance.

☐ The contract record states the applicable limits and exclusions, required insurance, and evidence of coverage.

☐ Each exception to the liability, indemnity, or insurance requirements identifies who approved it and records the decision about the remaining risk.

☐ Each category of data or artifact has a recorded approved purpose, recipient, and location. The record explicitly states whether use for product improvement or training is permitted.

☐ A contract review has verified that unapproved secondary use is prohibited and that requirements for authorization, enforcement, breach response, and evidence of deletion are stated.

☐ The retention schedule identifies each retained item, where it is held, its owner, retention period, and any legal hold.

☐ The retention schedule states how and when data will be returned or deleted and how completion will be verified. It explains how the obligation extends to subprocessors and states the remedy for failure.

☐ The exit terms define the triggers, authorizing role, and notice required. They specify support, service levels, export requirements, continuity, costs, and responsible owners.

☐ The exit completion record contains evidence that access and integrations have been removed, affected records corrected, and retention or deletion requirements met. Exit cannot be marked complete without this record.

☐ A dated transition exercise has tested suspension, replacement, or decommissioning of the service. It has tested the alternative service's capacity and how long it can operate, as well as export of data and system configurations, reconciliation of records and unfinished work, communications, and recovery time.

☐ The transition exercise has a named owner and backup. Its record gives the result, defects found, and retest date.

☐ The legal or regulatory owner records the jurisdiction, product and use classification, and applicable legal, professional, and hospital requirements. The assessment identifies required authorizations and reports, how modifications would affect them, and any findings or conditions.

☐ The dossier records the sourcing decision and its conditions, with a link from the registry.

E2.3 Supported configurations

☐ The configuration record states the approved maximum number of configuration groups for the service and the number currently in use. A configuration group contains variants of the same defined use. Evidence must support evaluating, monitoring, and supporting those variants together.

☐ The default is one configuration group per service. Any different limit has an approved rationale based on the service's capacity. This default is a proposed planning limit; evidence has not established it as a safety limit.

☐ Variants share one configuration group only when they are controlled and evidence establishes equivalence for the same defined use. The group record states the permitted differences in population, context, interfaces, data, permissions, versions, and workload and links the supporting evidence. Those differences remain within the group's approved operating limits.

☐ The hospital has confirmed that it can support each configuration group. Before the number of groups exceeds the approved limit, funded capacity and explicit authorization for a higher limit are secured, or an existing group is consolidated or retired. An exception cannot remove core controls.

E2.4 Hospital-developed components

☐ Hospital-developed components have an ongoing multidisciplinary team in place. It includes clinical and operational leadership, data science or machine-learning engineering, informatics, privacy, cybersecurity, legal, and quality or patient-safety expertise.

☐ Clinical and technical roles are assigned to named people, with backups and succession arrangements.

☐ The capacity assessment verifies that the build team can meet approved standards for development, testing, release, monitoring, incident response, and rollback or safe disabling. It documents the staff, systems, resources, and coverage supporting each duty.

☐ Funding and continuity arrangements are documented.

☐ Each capacity gap has an owner, due date, and recorded decision on whether the affected activity may proceed or must wait.

☐ The build inventory identifies every material component by version or repository and names its internal owner, backup, and release authority.

☐ The build inventory documents dependencies, support hours, funding, and retirement arrangements for each material component.

☐ For each signal of a material decline in service or performance, the internal responsibility matrix records the source and threshold. It names who detects, reviews, and responds to the signal, sets the required service level, and explains how an overdue response is escalated.

☐ For every retraining or update, the internal responsibility matrix names who requests, reviews, approves, and implements it. It specifies required evidence, funding, backup arrangements, and authority to approve rollback or safe disabling.

☐ Tests have confirmed that detection, review, and response are fast enough to prevent or limit the documented harm. For every safety-critical signal, the results compare the full delay from data collection through alert delivery to action with the time in which harm could occur. The measured delay includes the intervals between data collections.

☐ A test of the full monitoring and response sequence confirms that the expected signal is captured, its alert reaches the responsible recipient, and the required response follows.

☐ The test record states the result and assigns every unresolved gap to a responsible owner.

☐ Change procedures cover the model, data, code, prompts, thresholds, and infrastructure. They identify who may authorize a change, how versions are controlled, and the required impact screening. They also specify regression tests to confirm that previously accepted functions still work.

☐ The change procedure requires a record of approval and the evidence supporting release. It also specifies monitoring after release, triggers for rollback or safe disabling, and the handling of emergency changes.

☐ The handoff from research to operations identifies the exact version, specific use, and operating limits proposed for live operation. It records permission to conduct research separately from authority to operate the service.

☐ Operational owners have acknowledged unresolved findings and accepted their assigned actions.

☐ Every finding that blocks live use has been resolved through the required assurance and specialist reviews before production confirmation.

☐ The handoff record links the approved monitoring plan, confirmation that operational support responsibilities have been accepted, and the separate production confirmation.

☐ Validation, monitoring, incident response, release, rollback or safe disabling, and retirement each have a named owner and backup, with funding for the work.

☐ The capacity assessment compares expected workload and response targets with the build team's current capacity. It records gaps and the decision to proceed or hold.

E2.5 Purchased components and services

☐ The vendor intake record identifies the provider, exact product version and setup, intended use, and known limitations. It describes how data will move, where the service is hosted, and which subprocessors are involved.

☐ Vendor intake includes supporting evidence, privacy and security materials, and commitments for updates and incidents. It records support and exit arrangements, conflicts, unresolved questions, and the procurement decision.

☐ The supplier contact record names primary and backup contacts by issue type and states coverage hours.

☐ The supplier contact record specifies response and restoration times for each severity level and explains how to escalate an issue. It includes the latest contact-test result and the contractual remedy available if the supplier does not meet its obligations.

☐ The hospital and supplier have approved a map of observed service data flows, including every external recipient and storage location.

☐ Each difference between observed data flows and the approved map has a responsible owner, due date, and evidence confirming that it has been resolved.

☐ The contract establishes rights and custody for source and derived data, logs, prompts, templates, and outputs.

☐ The contract records hospital access, audit, correction, export, deletion, and transition rights and any approved exception.

☐ Hospital reviewers have recorded whether supplier evidence applies to the exact release and setup, population, workflow, and outputs proposed for local use.

☐ The hospital review identifies missing evidence and evidence that does not apply to the proposed use. For each gap, the record specifies the local evaluation or other assurance work needed to address it and the supplier's deadlines.

☐ The procurement record states whether the result supports selection, selection with conditions, further evaluation, or rejection.

☐ Each contractual performance claim identifies the release and configuration tested. It records the metric, population, data period, subgroup findings, comparator, acceptance threshold, and uncertainty.

☐ The contract specifies reporting deadlines and the hospital's rights to the data supporting each performance claim.

☐ For each contractual performance claim, the contract specifies when a decline in performance requires a response, the correction deadline, the remedy, and the hospital's right to verify that the response is adequate.

☐ Each model, hosting service, interface, and integration has a named hospital or supplier owner responsible for maintenance. The maintenance schedule specifies coverage, service levels, evidence delivery, change authority, dependencies, backup, funding, and escalation.

☐ For each signal of drift, failure, or performance change, the measure record names who detects it at the supplier and who reviews it independently at the hospital. It includes the measurement, data, comparison, uncertainty, and reporting fields required for material signals. Detection may be performed by a person or system. Drift means a departure from the conditions supporting approval, including changes in data, population, workflow, or system behaviour.

☐ For each signal, the monitoring plan sets the review and notification intervals. It names who responds, who corrects the problem, their deadlines, and who may pause the service. It states what evidence is needed to confirm that the issue has been resolved.

☐ The contract requires the supplier to monitor, investigate, share evidence, and take corrective action. It sets deadlines for corrective action and consequences for failing to meet these duties.

☐ The contract's update clauses specify which changes are covered, how much advance notice is required, how the exact release will be identified, the supporting evidence, and the time allowed for hospital review.

☐ The hospital retains the right to defer an update or require revalidation, and to receive emergency notice. It also retains the right to roll back or safely disable the affected service and activate an approved alternative, seek a remedy, or exit the arrangement.

☐ Incident clauses set deadlines for notice, acknowledgement, investigation, and correction according to severity.

☐ The supplier must preserve and provide logs and other evidence, and help identify and correct affected records.

☐ Incident arrangements name the hospital and supplier leads and specify support for regulatory work and communications. The hospital retains authority to pause the service.

☐ The incident arrangements define the contractual remedies and the evidence needed to confirm that an incident can be closed.

☐ The signed agreement implements the reviewed allocation of liability, indemnity, and insurance for patient or worker harm, privacy or cybersecurity events, intellectual-property claims, service failure, and regulatory non-compliance. It includes approved limits, exclusions, evidence of insurance coverage, and decisions about the risk that remains.

☐ The model provider is recorded as a subprocessor when it performs that role; otherwise it is recorded separately as a supplier or data recipient. The current subprocessor register identifies each organization's service, data access, purpose, location, retention, and further transfers. The contract requires advance notice of changes and preserves the hospital's right to object or exit.

☐ The register records evidence that deletion has been verified across all relevant subprocessors, with the latest verification date.

☐ For each named hospital role, the audit clauses specify which evidence it may obtain: logs, release records, evaluation and incident evidence, and corrective-action records. They also cover records showing where data and system components came from and how they changed.

☐ The audit agreement specifies delivery format and frequency, independent audit rights, and remedies.

☐ The date and result of the latest test of hospital access to the required evidence are recorded.

☐ Use of hospital data and outputs is limited to approved purposes.

☐ Use of hospital data or outputs for product improvement, training, or further purposes remains disabled unless separately authorized.

☐ The data-use record links the latest audit of whether use stayed within its authorization. It also links the evidence accepted to confirm completion of any breach-response action arising from that audit.

☐ The deletion record links the latest verification that deletion occurred and identifies any corrective action still outstanding.

☐ Contractual and technical evidence confirms compliance with the approved retention schedule, which identifies each retained item, its location, owner, retention period, and any legal hold. The evidence covers return or deletion methods and deadlines, obligations on subprocessors, the named verifier, confirmation of completion, and remedies.

☐ The signed exit schedule implements the agreed triggers, decision authority, notice, continuity, and export arrangements. It names who is responsible for removing access, correcting records, returning, retaining, or deleting data, managing costs, and testing that exit is complete.

☐ A transition exercise, supported by the supplier, has tested whether recovery can meet the approved target. It covers export of data, system configurations, and logs; continuity through a replacement or manual process; reconciliation of records and unfinished work; and communications. The results identify defects, responsible owners, and retest dates.

☐ The local assessment compares the exact release with approved criteria for the population, data path, workflow, users, capacity, and fallback. It records whether the release fits the proposed use.

☐ Every gap between the proposed service and local requirements has a recorded owner, mitigation, due date, and decision to accept, impose conditions, or reject.

☐ The hospital has accepted the design for independent local monitoring before governance approval.

☐ Before production confirmation, the hospital has verified that monitoring is in place and working. Readiness may be demonstrated in an authorized test environment without starting live use.

E2.6 Joint development and governance

☐ The partnership matrix assigns one accountable party and one named hospital owner to every decision and duty.

☐ The partnership record states which work and decision authority are delegated. It documents rights to evidence, service levels, funding, and backup arrangements, and names who resolves disputes, completes outstanding work, and verifies completion.

☐ The hospital retains authority to reject or pause the service.

☐ Both partners have approved a versioned data-flow map covering every data source, element, transfer, recipient, connection point, and storage or log location. The map and linked record cover exchanges between partners and with external recipients, hosting regions, suppliers and subprocessors, whether personal health information leaves the hospital, authority for access, and any permitted further use. Retention and deletion duties are assigned.

☐ A dated comparison of the approved data-flow map with observed flows records the findings and the action taken on each difference.

☐ The partnership agreement names rights holders and hospital custodians for data, logs, prompts, templates, outputs, and jointly developed artifacts.

☐ The partnership agreement records access, change, licensing, audit, export, deletion, intellectual-property, dispute, and exit rights.

☐ The partnership matrix assigns responsibility for validation, monitoring, updates, maintenance, incident response, correction, rollback or safe disabling, communications, and retirement for each component. It identifies who may authorize changes and releases, the independent review required, evidence requirements, and deadlines. The hospital retains its decision rights.

☐ The partners' assurance agreement identifies the exact release covered. It specifies agreed measures and thresholds, rules for assessing the population and subgroups, evaluation methods, access to evidence, reviewers, and review interval. It also names who resolves disputes.

☐ Each joint update has a record of the proposed change, who proposed it, and which party controls it. The record documents the impact assessment, evidence, and acceptance criteria, followed by approval, release, monitoring after release, arrangements for rollback or safe disabling, and resolution of any disagreement.

☐ The partnership incident schedule sets notification and response deadlines for each party and specifies what each party must do with evidence. It names the investigation and communication leads, who may authorize containment, and the corrective-action owner. It explains how disputes are escalated and what evidence is needed to close the incident.

☐ The partnership agreement implements the allocation of liability, indemnity, and insurance for patient or worker harm, privacy or cybersecurity events, intellectual-property claims, service failure, and regulatory non-compliance. It includes approved limits, exclusions, evidence of insurance coverage, and decisions about the risk that remains.

☐ Each partner maintains a current subprocessor register identifying each organization's service, data access, purpose, location, retention, and further transfers. Other external suppliers and data recipients are recorded separately. The agreement requires advance notice of changes, preserves the hospital's right to object or exit, and covers deletion obligations across all relevant subprocessors and verification dates.

☐ The agreement names the party responsible for producing each required log, release record, evaluation record, incident record, and corrective-action record. It also assigns responsibility for records showing where data and system components came from and how they changed. It specifies the format, frequency, and deadlines for hospital access, independent audit rights, tests of access, and remedies for non-delivery.

☐ For each category of data or artifact, the partnership agreement identifies the permitted purposes, parties, and locations. The data-use schedule records prohibited uses, any permission for training or product improvement, and the authorization for permitted uses. It also specifies audit checks, breach response, and evidence of deletion.

☐ The joint retention schedule identifies every retained artifact, its responsible party, the system and location holding it, and its retention period or hold.

☐ The schedule states the return or deletion method and deadline and extends the obligation to subprocessors.

☐ A named verifier has recorded the evidence and date from verification of return or deletion, along with the remedy for non-compliance.

☐ Exit terms define measurable triggers, decision authority, and notice requirements.

☐ The continuity plan identifies what must continue during transition and how data, models, configurations, and logs will be transferred.

☐ The agreement assigns responsibility for removing credentials and integrations, correcting records, retaining or deleting data, and communicating with stakeholders.

☐ Exit arrangements state the cost, responsible parties, deadlines, and test used to confirm completion.

☐ Before production confirmation and annually thereafter, the partners run an exercise to confirm that authorized staff can suspend the service and maintain continuity. It includes transferring data, models, configurations, and logs.

☐ The exercise tests removal of credentials and integrations, the capacity of manual or replacement services, communications, and whether the work can be completed within the approved time.

☐ The exercise results identify corrective actions, responsible owners, and retest dates.

E3 Local evaluation and permission to start

E3.1 Validation and human factors

☐ Before reviewers see results, the validation plan specifies the release and setup, population, evaluation period, sampling approach, comparator, workflow, and users it covers.

☐ Before testing begins, the validation plan sets the measures, minimum acceptable subgroup results, decision rules, and treatment of uncertainty. It names the reviewer and approving authority.

☐ The validation report identifies the operating setup tested for live use, including its data flow and integrations, and the reference standard used to assess results. It compares results with the acceptance criteria and reports uncertainty, deviations, limitations, corrective actions, and the final decision.

☐ Representative intended users have completed scenarios covering ordinary and difficult work, time pressure, disagreement, errors, escalation, and fallback.

☐ The human-factors assessment compares observed errors, delays, over-reliance, workarounds, and accessibility barriers with the criteria set before testing. The results are recorded.

☐ Every failed criterion has a recorded corrective action and owner, and the final decision is retained.

☐ The local assurance assessment identifies the relevant equity questions, groups, and measures. It reports uncertainty and missing data and applies the decision rules set before testing. It records mitigation, the responsible owner, and the result.

E3.2 Approval and implementation readiness

☐ The monitoring plan is approved before production confirmation.

☐ Authorized testing and readiness checks verify that monitoring feeds and calculations work. They confirm that monitoring will operate from the start of live use.

☐ The monitoring plan records the measures, baselines and thresholds, who reviews each measure, how often, and the response deadlines.

☐ The monitoring-operations record states where reports can be found, how concerns are escalated, and what evidence is required before a stopped service may restart.

☐ The dossier index records the registry link, dossier location and version, local form identifier, and dossier coordinator. It links the current baseline, factsheet, evidence, local assurance, data map, hazard assessment, and all supporting lifecycle records, including user instructions, monitoring, competency, transparency, contracts, changes, incidents, and retirement.

☐ The governance decision identifies the release, configuration group, permitted and prohibited uses and operating limits, risk tier, remaining risk, and supporting evidence. It states which preparations may begin. Before live use, a separate production confirmation must verify that the actual system meets the conditions of approval and is ready to operate.

☐ The governance decision names the approver and records every condition, its responsible owner, deadline, and who may accept evidence that it has been met. It includes monitoring and stop rules, the approval's expiry, and the next review date.

☐ Before release, production confirmation verifies every prerequisite for live use.

☐ Before receiving access, each Tier 3 or Tier 4 user has demonstrated that they can stay within the approved use, perform required verification, recognize prohibited uses, and use fallback. The assessment applies pass criteria specific to the user's role.

☐ The competency assessment also covers escalation and incident reporting.

☐ The competency record gives the assessment result, decision on access, and events that require reassessment.

☐ The early-adopter plan identifies the initial users and the limits of the approved use. It specifies which user practices will be observed and checked, the early review date, and the criteria for expansion.

☐ A completed exercise for the proposed production release has tested rollback, safe disablement, or restoration. It has assessed recovery time, reconciliation of records, continuity through a manual process, and communications. The exercise record identifies the decision authority, date, result, and any unresolved defects.

☐ The disclosure record states whether disclosure is required and identifies the authority for the decision.

☐ Where disclosure is required, the approved approach specifies what people will be told, when, and in which accessible format. It provides a practical alternative where feasible or required, and routes for questions and corrections. Any limits on alternatives and how concerns will be addressed are recorded.

☐ The disclosure record names the responsible owner, effective date, and evidence that the approved approach has been implemented.

E3.3 Expansion

☐ The expansion request describes the proposed expansion, identifies every change, and states which parts of the approved baseline would stay the same. It explains how changes will be assessed together when they are combined.

☐ Reviewers consider the expansion evidence for safety, performance, equity, workload, capacity, value, patient and user experience, adherence to required practices, and incidents. For any omitted area, the reviewer accepts a recorded explanation of why it does not apply.

☐ For each applicable area of expansion review, the assessment records the observation period, denominator (the total quantity used to calculate a rate), acceptance criteria, and result. It records uncertainty, missing evidence, open incidents or actions, the responsible owner, and decision.

☐ Reviewers have compared the service receiving the expansion with the approved use, covering site, population, language, volume, features, integrations, actions, and the system's freedom to act independently. They have verified both the proposed differences and claims that an aspect will remain unchanged.

☐ For each aspect of expansion, the comparison records the approved baseline and proposed difference, with the evidence, risk or capacity gap, mitigation, responsible owner, and decision.

☐ The expansion assessment covers clinical operations, technical support, monitoring, training, incident response, continuity, and rollback or safe disabling for the proposed wider use. Each function has an owner and a measurable capacity standard.

☐ Current evidence establishes whether each function meets that standard. Where it does not, the record identifies which expansion activity depends on it and must remain on hold.

☐ An unmet capacity standard blocks the part of expansion that depends on it. A narrower scope may be authorized only if it meets the applicable standards.

☐ The expansion decision records who authorized it, its conditions, expiry, and first review date. It clearly identifies any additional use or operating scope that remains unauthorized.

☐ Before the expansion goes live, production confirmation has verified that the actual system meets the approval for the expanded use.

E4 Monitoring, change control, and retirement

E4.1 Monitoring operations

☐ Retained test results demonstrate that the full monitoring process, including its data feeds and calculations, meets the approved criteria.

☐ The monitoring operations record identifies the monitoring report, accountable owner, and backup. It sets targets for data completeness and reporting delay, states the review schedule, and explains how an overdue review is escalated. It also identifies the fallback to use if monitoring fails.

☐ Monitoring covers every material risk. Where applicable, it covers safety, performance, data quality and drift, subgroup effects, use and workflow, value, reliability, and incidents. A reviewer accepts the reason for omitting any area.

☐ Each measure's record states what is measured, its source, denominator, calculation version, baseline, comparison period, and current result. It specifies minimum data requirements, uncertainty and how it will be handled, required site or subgroup reporting, and the response to missing or late data. It names the reviewer and backup, review and reporting intervals, report recipients, last and next review dates, overdue-review response, and the required response at each threshold.

☐ The risk assessment considers every failure pattern: target and representation bias; data and label noise; covariate shift and out-of-distribution inputs; prevalence changes and label shift; concept drift and changes in the target definition; data leakage and shortcut learning; prediction and calibration errors; data pipeline and integration failures; hallucination and other generation or retrieval errors; agent goal and execution failures; reward hacking and specification gaming; security and privacy failures; automation bias and other human or workflow failures; and failures in monitoring or the response to problems. The risk assessment and monitoring record link the applicable hazards, safeguards, monitoring measures, and gaps in detection. A reviewer has accepted and recorded the reason for any exclusion.

☐ The review plan includes random case review, targeted review, and routes for patient and staff concerns to identify effects not yet in the hazard list. Reviews include relevant cases that produced no alert and interactions that failed. The plan names reviewers and sets review intervals and escalation deadlines.

☐ An authorized exercise verifies critical monitoring calculations, completeness of source data, notification, response, and fallback. Its record identifies the evidence, faults, owners responsible for correction, and retest dates. The record states any uncertainty caused by too few observations to support a conclusion.

☐ Each threshold has a response record specifying the trigger, required action, authorized primary and backup responders, deadline, and escalation route. It states how uncertainty, insufficient data, or reporting delay affects the decision, and specifies fallback, communications, and any review of affected records. Any temporary emergency arrangement has separate authorization, a time limit, and equivalent protection; mandatory duties still apply. Reaching a stop threshold requires immediate suspension of the affected function, activation of fallback, and incident escalation. Continued use under an approved restriction or restart requires the specified evidence. Restart also requires authorization for the resumed scope and a separate production confirmation that the actual release and scope meet the approval.

E4.2 Incident readiness

☐ A dated exercise demonstrates that users can access the incident-reporting route and assign the correct severity.

☐ The incident-exercise record compares the results with approved targets for notifying the on-call team, preserving evidence, containment or pausing the service, and issuing required notices.

☐ Each defect found in the exercise has an owner, deadline, and evidence confirming that it has been resolved.

E4.3 Reassessment and change

☐ The dossier records the tier's maximum review interval.

☐ Threshold breaches, incidents, material changes, supplier notices, expansion, new or changed evidence, or overdue corrective actions trigger an earlier review.

☐ Each material change has been logged, assessed, and authorized before release. Its record includes the baseline, supporting evidence, results from tests of previously accepted functions and safeguards, conditions, and arrangements for rollback or safe disabling.

☐ Before a changed release enters live use, production confirmation has verified that it is ready and meets the conditions of approval.

☐ Emergency changes have followed the authorized process for containing harm and have been reviewed within the required time after the change.

☐ The update procedure identifies who may authorize a change. It requires an impact assessment before release, evidence from tests of previously accepted functions and safeguards, approval, version control, and communication.

☐ The update procedure states how the system will be monitored after a change, who is responsible for rollback or safe disabling, and the deadline for reviewing an emergency change.

E4.4 Retirement

☐ Measurable retirement triggers cover value, safety, performance, equity, support, suppliers, and loss of monitoring.

☐ The retirement decision names its owner, records its approval, and sets the deadline and effective date. It includes the continuity plan and requirements for confirming completion.

☐ Verification confirms that access, credentials, integrations, automation, and scheduled jobs are disabled.

☐ Verification confirms that open work and affected records have been reconciled.

☐ Retirement is marked complete only after verification that retention controls are in place, required export and deletion are finished, and the replacement service or workflow is operating as approved.

☐ Each remaining retirement action or lesson has an owner and date.

Appendix F Checklists for system capabilities

These checks apply according to what the system will do in the proposed use. Prediction checks cover systems that predict outcomes, assign categories, detect findings, or rank cases. Generation and retrieval checks cover systems that create content or retrieve information; scribe workflow checks also apply whenever the system captures information from an encounter or prepares documentation. Action and memory checks apply whenever the system can take actions or retain information for later use, regardless of its product name. A product may need checks for several capabilities. These checks supplement the risk-tier and lifecycle requirements. An item is complete only when the evidence required for the current decision has been accepted and the verifier and date are recorded. Evidence due only at a later decision may remain pending while work that does not depend on it proceeds. A reviewer must accept the reason for marking any function-specific item N/A; N/A cannot waive a mandatory control.

The functions enabled for the proposed use determine which checks apply; the consequences of its proposed use determine the risk tier. For example, a system that generates content without retrieving information may have no retrieval source to assess. A product sold as a scribe still needs the action controls if it can act on its output. The reviewer records the required sections and explains any individual item that does not apply. A configuration group brings together different setups for the same defined use. Evidence must show that they can be evaluated, monitored, and supported together within approved operating limits.

F1 Prediction, classification, detection, and prioritization

F1.1 Approved use and evidence

☐ Each output has a record of the decision or action it may influence, the cutoff or ranking used to guide that decision, and the possible consequences.

☐ The output record names the intended user and states what information they see, the extent to which they may rely on the output, and how quickly they must respond. It identifies where approval is required, how the user can override the output, and the practical alternative available.

☐ The baseline workflow and the non-AI process used for comparison are described, including steps, roles, volume and timing, errors and near misses, workload, and access.

☐ The comparison record identifies the measures of outcomes and equity, the period covered by the data, and known limitations.

☐ The validation plan identifies the exact version and configuration, decision threshold, population, evaluation period, cases, intended users, and workflow. It describes how local data will reach and be processed by the system, the comparator, and the reference standard used to judge correctness. Before testing, applicable criteria for benefit, performance, serious errors, subgroups, and equity are set, with uncertainty rules and Pass, Fail, or Insufficient evidence decisions. Critical criteria are identified explicitly. Every applicable critical criterion must Pass; N/A requires the reviewer's acceptance that the criterion does not apply.

☐ The validation report presents the results, uncertainty, limitations, and departures from the test plan. It records the resulting decision, reviewer, date, and approver.

☐ Data leakage and shortcut learning risks have been assessed. The relevant tests are recorded, including separation of training and evaluation data and performance under changes to suspected incidental cues.

☐ For risk predictions, the plots or measures used to assess calibration—how closely predicted risks match observed outcomes—are specified before analysis.

☐ For risk predictions, the validation report states how closely predicted risks match observed outcomes overall, at required sites, and in required subgroups.

☐ Calibration is assessed across the risk ranges that matter clinically, with acceptance criteria and uncertainty reported.

☐ If predicted risks are adjusted to better match observed outcomes (recalibration), the method and resulting decision are recorded.

☐ The equity assessment identifies relevant groups and how the system could affect them. It reports each group's sample or event counts, measures, uncertainty, and missing data so reviewers can assess the strength of the evidence.

☐ The equity record documents how findings for each subgroup were judged against the decision rules. It records the results, actions to reduce risk, responsible owner, and final decision.

F1.2 Human factors and monitoring

☐ Intended users have passed scenarios involving incorrect, missing, delayed, and conflicting outputs and large volumes of outputs. The scenario results are retained.

☐ The scenarios also test disagreement with the system, escalation, and downtime.

☐ Scenario results meet the criteria set in advance for response time, overrides, demands on users, workarounds, and failures.

☐ Every initial scenario failure has a documented corrective action and decision on retesting.

☐ The monitoring plan includes applicable measures of data quality, performance, calibration, use, workflow, outcomes, equity, and incidents. Any omitted domain has a recorded explanation.

☐ Each monitoring measure has a recorded definition, source, calculation version, baseline, comparison period, and denominator specifying the population or events counted. Its record states the minimum data needed, how uncertainty is handled, required site or subgroup reporting, the response to missing or late data, and review and reporting intervals. The reviewer, backup, report recipient, last and next review dates, escalation for overdue review, and link from the current result to its required response are recorded.

☐ The response record specifies the response required at each threshold, including any rule for recalibration or retraining. Meeting a stop criterion requires immediate suspension of the affected function, activation of fallback, and incident escalation. Restart requires accepted evidence and production confirmation.

F1.3 Change and continuity

☐ Retraining and changes to system features, data sources, settings, populations, thresholds, workflows, or suppliers are assessed against the approved definition of a material change. Each material change has a versioned change record.

☐ The change record names the reviewer and decision authority and records the date, conditions, release, arrangements for rollback or safe disabling, and monitoring results after the change.

☐ Before production confirmation, an exercise demonstrates that the non-AI alternative can handle the expected peak workload. Results are assessed against approved criteria for downtime, safety, staffing, communications, checking and resolving differences in records or work, and restoring service.

☐ The exercise results and defects are recorded, with assigned corrective-action owners and a retest date.

F2 Generation, retrieval, and ambient documentation

F2.1 Approved use and evidence

☐ The approved-use record identifies the users, sites, population, setting, shifts, exclusions, expected volume, and maximum approved scale, together with inputs, outputs, and place in the workflow. It states the extent to which users may rely on the system. Prohibited data, content, actions, and later uses of the output are listed separately.

☐ The generative-AI data rules specify which data elements and sources the system may or may not use and which recipients and subprocessors may or may not receive them. A subprocessor is an organization engaged by a service provider to process data on its behalf as part of the hospital service.

☐ The data rules state where the service is hosted and which jurisdiction applies. They record access and logging rules, retention and deletion periods, any permitted uses beyond the original purpose, and controls on personal health information (PHI) leaving the hospital.

☐ The data rules explain the approved routes for requesting a correction and reporting an incident.

☐ Every prompt, template, and guardrail used for a Tier 3 or Tier 4 use has a version-controlled record. Prompts instruct the system, templates structure its content, and guardrails are controls intended to constrain its behaviour.

☐ The configuration record names the person responsible for the setup and who may approve changes. It records the approved use, configuration group, test cases, criteria for passing, and release date, and specifies the safe setup or service state to use if a change must be reversed or the service disabled.

☐ Before release, a comparison confirms that the actual production system matches the approved configuration.

F2.2 Human factors and monitoring

☐ The output-review record names who checks each output requiring review. It states what they must verify and when they must correct and formally confirm the output (attestation). Release without prior review is permitted only for a separately approved use and within explicit limits on what the system may do.

☐ Reviewers have the authority and manageable workload needed to check outputs.

☐ The output-review controls state what reviewers must complete and how completion can be observed. Any minimum review time is justified by the task.

☐ The output-review procedure specifies how to handle disagreement and exceptions, which shortcuts are prohibited, and how compliance is audited.

☐ Before testing, an agreed table of test cases covers the exact setup proposed for live use, intended users, populations, languages, workflows, and sensitive content. It identifies the ways the system could fail that each test must examine.

☐ Acceptance criteria are fixed before evaluation begins, with errors weighted according to severity.

☐ The assurance report presents test findings, departures from the plan, conditions not tested, and the resulting decision.

☐ The monitoring plan defines the types of unsafe output whose frequency will be monitored, signs of use outside the approved scope, and the required review of affected records. Each measure has a recorded source, calculation, denominator, baseline, comparison period, minimum-data and uncertainty rules, and required site or subgroup reporting. Reviewers, backups, review and reporting intervals, recipients, last and next review dates, and responses to missing data or overdue review are specified. Each threshold has a required response, authorized responders, supplier duties, response and escalation deadlines, and evidence needed for continuation or restart.

F2.3 Change and continuity

☐ Contract terms state how far in advance changes must be reported, how each release will be identified, what evidence must accompany a change, and the rules governing which features may be enabled by default.

☐ The hospital's change procedure states how to assess a change's impacts and which aspects of the system and its use must be validated again. It identifies who approves the change and how approval is obtained. It covers deferring or disabling a change, restoring a previous approved setup, handling emergencies, and informing users.

☐ For patient communications, documentation, and summaries, the hospital records whether AI use must be disclosed and the authority supporting that decision.

☐ Where disclosure is required, the approved approach identifies the audience, timing, and accessible means of communication. It provides a practical alternative where feasible or required, and routes for questions and corrections. Any limits on alternatives and how concerns will be addressed are recorded.

☐ The disclosure record names the responsible owner and records evidence confirming that the approved approach has been implemented.

F2.4 AI scribe workflow checks

These checks apply whenever the system captures information from a clinical encounter (ambient capture) or prepares clinical documentation. Patients must receive the required information and opportunity to choose before capture starts; later approval cannot authorize an earlier recording. The required assurance and release checks are complete before live use, with evidence linked to the exact setup authorized for use.

☐ A versioned record of the scribe's approved settings and workflow describes how it captures information, which sources supply background information, and what it produces. It states how the output enters the record and the authority permitting any further activity based on that output.

☐ The configuration record documents data handling and retention, the versions of the model, prompt, template, and guardrails, and the approved clinical context.

☐ Patients receive clear information in the approved languages and formats before capture begins.

☐ Start and stop controls are visible.

☐ The capture procedure explains how to handle people whose voices or information may be captured although they are not participating.

☐ Procedures for retaining information, correcting errors, handling complaints, and documenting care without AI are defined and tested.

☐ The scribe test plan specifies the exact configuration and source of test cases. It justifies the sample size and coverage of intended specialties, languages, and difficult cases.

☐ Before testing, the assurance plan classifies error severity and sets the rules for Pass, Fail, or Insufficient evidence.

☐ The assurance plan separately defines when an evaluation or live operation must stop.

☐ The scribe test report identifies omissions and hallucinations (fabricated or unsupported content) in the output. Where applicable, it reports errors in who said what, the order of events, medication, allergies, plans, and follow-up. It states the uncertainty in the findings and the resulting decision.

☐ Controls keep outputs as unsigned drafts until the required review and formal confirmation are complete. Automatic finalization is prohibited unless authorized as part of a separately governed use.

☐ Enabling coding, patient messages, referrals, orders, searches across the chart, completion of structured fields, or task execution requires review and authorization of that function as a separate or materially changed use.

☐ Every enabled action-taking feature meets the applicable requirements for permitted and prohibited actions, technical permission limits, human authorization and independent shutdown, auditability and retained information, and testing, monitoring, incident response, change, and continuity controls. Evidence that these requirements are met is accepted before the decision or activity it supports.

☐ The scribe record links the approved configuration and disabled features to the relevant supplier change notice.

☐ Test results supporting the proposed release are retained, including regression-test results showing whether previously accepted functions and safeguards still work.

☐ The scribe record assigns responsibility and deadlines for scribe audits, monitoring, correction, rollback or safe disabling, restart, and exit.

☐ The scribe stop procedure requires immediate suspension of the affected function, activation of fallback, and incident escalation when a stop threshold is reached.

F3 Agentic AI and persistent memory

F3.1 Action boundaries

☐ The action register states what triggers each permitted action and what must be true before it runs. It identifies where the action takes effect, how it is linked to the correct patient or context, its limits, and the required final state when it finishes. It records the required human approval or separately approved permission to act without prior confirmation.

☐ The action register lists prohibited actions and the conditions in which an action is prohibited separately.

☐ Technical tests confirm that each action uses only approved credentials and permissions, stays within its limits, stops when its authorization expires, and reaches its required final state.

☐ An approved list of tools, software connections (APIs and other integrations), actions, access credentials, and permissions is maintained and enforced. An API lets systems exchange data or carry out actions. Permissions specify which data a connection may access and what it may do. Technical controls block anything outside the approved list.

☐ Each approved tool, API, integration, action, credential, or permission has a record of its owner, prerequisites for use, limits, required approvals, expiry, and audit result.

☐ The dossier describes how independently the agent can act, what effects it could have, how complex its goals are, and the range of tasks it can perform. It explains how the consequence tier reflects the permitted actions, affected systems and populations, time operating unattended, required human approvals, and controls for stopping or reversing actions. [57]

☐ The agent description has been reassessed after any change to its goals, tools, or permissions.

F3.2 Human authorization and shutdown

☐ For each action requiring prior human approval, the review record identifies when the reviewer must decide and names the primary and alternate reviewers. It states what evidence they see, how much time they have, what they may edit, when approval expires, and how to escalate a concern.

☐ Technical testing confirms that the agent acts only within its approved authority and cannot bypass required human approval.

☐ Testing confirms that a named operator can independently disable the system's ability to write data or restrict the service within the approved maximum time.

☐ The stop procedure requires that disablement or restriction begin immediately at a stop threshold and that fallback be activated.

☐ The continuity record identifies alternatives that allow staff to view information without changing it (read-only access) or continue work manually. It states who may activate and use them, their tested workload capacity, and how long they can operate, with links to the communication plan and latest exercise result.

☐ Every scheduled shutdown test and every test triggered by a change has a record of the date, release and permissions tested, time to disable the service, recovery, defects, responsible owners, and retest status.

F3.3 Auditability and memory

☐ An audit record preserves the agent's goal, governing instructions or policy, information sources, proposed plan, approvals, actions, errors, and final state, including changes in its task status, progress, and stored working information.

☐ A dated test using a sampled transaction confirms that authorized reviewers can reconstruct what happened from the recorded inputs, outputs, actions, changes in state, and system version. The test need not reproduce an identical output when outputs can vary between runs.

☐ The assurance record identifies audit evidence the hospital cannot obtain and the risk this creates. The affected activity remains on hold if reviewers cannot complete the required checks without that evidence.

☐ Any accepted alternative control for unavailable audit evidence has a named owner and evidence that it works.

☐ Verification confirms that audit logs capture the required events, identities, patient or context, timestamps, version, and result.

☐ The logging plan states how long records are kept, who can access them, and how they are protected from alteration or loss. It sets review and alert intervals and deadlines for providing the records as evidence. Verification confirms that the agent cannot alter past audit records or its evaluation criteria.

☐ The result of the log retrieval test is recorded.

☐ The state-and-memory record defines what the system may retain about the current task (state) and remember for later interactions (memory). It specifies the permitted fields and sources, how retained information may influence later activity, and the required separation between patients and users.

☐ The state-and-memory record names the owners responsible for access, correction, retention, deletion, and audit of state and memory.

☐ Tests find no prohibited transfer of stored task information or memory between patients, users, or contexts in the scenarios examined. Their coverage and remaining uncertainty are recorded.

☐ The latest test showing that disabling or resetting the system clears stored task information as intended is retained.

☐ Monitoring over time checks for harmful reinforcement, inappropriate dependence on the system, and continued activity that no longer fits the patient's changing goals. Harmful reinforcement includes repeated responses that strengthen a harmful belief or behaviour.

☐ Tests confirm that a person can take over, errors can be corrected, and memory can be reset as required when uncertainty or risk increases.

F3.4 Testing, monitoring, and incidents

☐ Before testing, the test plan identifies the agent, tools, permissions, and software connections that match the proposed live service. It lists the test scenarios and criteria for safety and recovery.

☐ The agent test report records departures from the plan, unresolved failures, corrective actions, and regression-test results showing whether previously accepted functions still work. It records the resulting decision.

☐ For Tier 3 or Tier 4 agents, simulations use permissions representative of live use while preventing real effects outside the test. They test normal and adversarial use, including prompt injection (attempts to redirect the system through unauthorized instructions in inputs or retrieved content) and other deliberate attempts to defeat safeguards. They also test failures, recovery, disabling actions that change data or other systems, and fallback.

☐ The simulation record states the results and the decision to proceed or hold.

☐ Tests assess whether the agent stops or requests authorized help when it cannot complete a task within its permitted access and actions.

☐ Claimed task completion has been compared with evidence from the receiving system. Any mismatch and its effect on the test decision are recorded.

☐ The agent monitoring plan measures unauthorized or failed actions, changes to permissions, events that prevent the system from writing or changing data, partly completed transactions, breaches involving retained information or context, repeated attempts, and use of alternative workflows.

☐ Each agent monitoring measure has a recorded definition, source, calculation, denominator, baseline, comparison period, minimum-data and uncertainty rules, and required site or subgroup reporting. Its record names reviewers and backups, review and reporting intervals, recipients, last and next review dates, and responses to missing data or overdue review. Thresholds are set for each applicable action, with required responses, authorized responders, deadlines, escalation, and evidence required for restart. Reaching a stop threshold requires immediate suspension of the affected function, fallback, and incident escalation; restart requires accepted evidence, authorization for its scope, and production confirmation.

☐ The agent incident record reconstructs attempted and completed actions, the permissions used, affected records and systems, unfinished actions, and any reversals or other corrections.

☐ The incident record documents how harm was contained, how the system was prevented from writing or changing data, what was corrected, and who was notified. It specifies the evidence needed for restart and the escalation deadlines.

☐ Each change to tools, connectors, memory, coordination of actions, or approval policy has a version and a documented assessment of its effects.

☐ For each material change, access permissions and data flows have been reviewed. Relevant workflow tests and attempts to defeat safeguards have been repeated as regression tests: checks of whether previously accepted protections still work.

☐ Each material change has recorded authorization.

☐ Each material change has recorded production confirmation.

☐ The changed release is monitored. The hospital retains the ability to roll it back or to safely disable it and use an approved alternative.

☐ A dated exercise tests whether the system can independently be prevented from changing data or other systems while read-only access or manual work continues, and whether queued and partly completed actions can be accounted for and resolved. Communications, recovery results, decision authority, defects, and retest dates are recorded.

Appendix G Procurement and external data flows

Procurement and external-data review is complete before the hospital commits to a supplier or a route for data transfer. The review establishes whether the hospital will have the evidence, access, and controls needed to run the proposed service. Supplier checks cover components obtained from suppliers or partners and the external dependencies involved. An internally developed component does not acquire supplier requirements solely because another component is purchased. The sourcing record links each finding to the decision it supports, with the reviewer and date. The review is repeated after a material change to the contract, system design, or supplier. Local assurance is still required before operational use.

A vendor proof of concept (POC) evaluates a supplier's product or setup. It must be time-limited and occur before purchase or commitment to a particular setup. A POC is required when existing evidence or an equivalent authorized evaluation leaves important questions unanswered about suitability, performance, connection to hospital systems, or fit with the workflow. If a separate POC is omitted, the procurement owner and responsible governance authority record the evidence and reasons, subject to local policy. An internal build without vendor selection may mark this item N/A, but still needs local assurance. POC authorization permits only the evaluation: access to data or live systems needs separate clearance, and outputs cannot be relied on in care. All other applicable approvals remain required. Any live pilot needs the relevant limited operational approval and production confirmation, whatever it is called.

G1 Procurement decision

G1.1 Proposed use and supporting evidence

☐ The procurement proposal defines the intended use and maximum proposed operating limits. It identifies who uses the system and why, its inputs and outputs or actions, who receives them, and the affected workflow. It records the level of AI involvement for each task, human decision authority, and the status of resulting records, such as draft or final. The limits specify permitted users, sites, settings, shifts, populations, exclusions, expected volume, and maximum scale.

☐ The proposal identifies who may not use the system, prohibited data and settings, prohibited features and actions, and prohibited ways of relying on its output.

☐ Supplier evidence identifies the exact version tested, datasets, population, evaluation period, what it was compared with, and measures used.

☐ Hospital reviewers have checked that reported results include uncertainty, serious errors, subgroup findings, and exclusions and limitations relevant to the proposed use.

☐ Hospital reviewers record where supplier evidence does not apply to the local use and the resulting procurement decision.

☐ The proposal or contract states which system setup, mapping of data between systems, and test environment the supplier must support for local assurance.

☐ The proposal or contract specifies the documentation, logs, representative cases or interfaces, and supplier personnel needed for local assurance.

☐ The agreement makes responsibilities for response times, evidence-delivery dates, and help with correcting problems enforceable.

G1.2 Cybersecurity and contract requirements

☐ The supplier has provided a dated map of the exact service's components, connections, and data flows. The map identifies every external recipient and location, how long information is retained, and how it can be deleted.

☐ Hospital reviewers record missing or inadequate evidence about the system design or data flows and the decision on each gap.

☐ Privacy and cybersecurity reviewers assess the evidence against the system's actual components and connections. They record the applicable requirements, findings, conditions, and remaining risk.

☐ Each privacy or cybersecurity gap has a recorded owner, due date, and evidence confirming resolution.

☐ The contract specifies which changes require notice, the advance-notice period, release identifiers, and accompanying evidence. It includes a route for emergency notice.

☐ The hospital retains authority to assess the effects of changes, require revalidation, approve or defer a change, and disable or roll back the service.

☐ The contract specifies the remedy if the supplier fails to meet an obligation concerning changes.

☐ The contract records required logs, retention periods, hospital access rights, and delivery deadlines.

☐ Hospital reviewers have verified that the hospital can obtain affected data and other evidence and access relevant supplier personnel for audits and investigations.

☐ The incident clauses assign responsibility and deadlines for preserving evidence and correcting problems. They state how unresolved issues are escalated and what remedies apply if obligations are not met.

☐ Secondary use, model training, product improvement, or resale of hospital data is prohibited without explicit authorization.

G1.3 Continuity and accountability

☐ The exit clauses state the conditions under which the arrangement may be ended, the required notice, and the time allowed for transition. They include the service commitments, assistance, and costs needed to keep work going during the transition.

☐ Those clauses also cover exporting data and system configuration, evidence of deletion, removal of system connections and access credentials, responsibility for alternative workflows, and remedies if obligations are not met.

☐ Contract reviewers identify the clauses assigning responsibility for clinical harm, inaccurate output, privacy and security failures, downtime, regulatory failures, and supplier failures.

☐ The contract review records insurance and indemnity arrangements, exclusions, unresolved accountability gaps, and approval authority.

☐ The contract gives the hospital rights to obtain evidence, investigate problems, restrict or pause the service, export or delete the covered data and other artifacts, and receive support when ending the arrangement.

☐ A technical exercise or discussion-based rehearsal has tested the exit plan. It covers disabling the service, continuing work manually, communications, export and deletion, and who performs each exit duty.

☐ Every issue left unresolved by contract review or the continuity exercise has an assigned owner and deadline. The procurement decision states how each issue will be handled.

G2 Comparative vendor scoring

Vendor comparison assesses how well products meet the same defined need, using one scorecard structure with separate sections where the approved criteria require them. The comparison records unavoidable product differences and comparisons that cannot be made fairly. It explains how the evidence supports the final selection, including uncertainty and any approved departure from the scoring method.

☐ Vendor evaluations use the same defined need, scenarios, comparator, evaluators, and core measures. Unavoidable product differences are documented.

☐ The rules for each part of the scorecard are approved before evaluation: the scoring scale and its weight in the total, required evidence, evaluator, minimum acceptable result, and any condition that means the proposal cannot pass.

☐ Every vendor score is linked to supporting evidence. The assessment records missing information, disagreement between evaluators, uncertainty, assumptions about total cost, and the resulting judgment.

☐ Recorded decisions to impose conditions or reject a proposal follow the scoring rule approved before evaluation.

☐ The selection record identifies funding and in-kind support, such as donated products, staff time, or services, before vendor selection.

☐ Evaluator and decision-maker conflicts are declared and recorded before selection. The record explains how they were managed and identifies any withdrawal from evaluation or a decision because of a conflict.

☐ Any departure from the agreed scoring method has been approved.

☐ The final vendor comparison explains how the evidence and remaining uncertainty support the decision.

G3 External data sharing and transfers

☐ The external-data review maps every external transfer and recipient using the shared data and integration records. These identify data sources, sensitivity, patient-data status and form, authorized purpose, quality and lineage owners, known limitations, connections and data formats, user and software identities and permissions, release responsibility, possible failures, and monitoring.

☐ When initial screening requires further review, named privacy and cybersecurity reviewers assess the actual system design, data flows, threats, and safeguards.

☐ The external-data review records the basis for permitting the proposed transfer, the notices required, supporting supplier evidence, and remaining risk.

☐ The external-data review records its conditions, approval decision, and any expiry.

☐ Each condition has an owner and due date, with a link to the evidence accepted as confirmation that it has been satisfied.

☐ Each mapped data element has a recorded purpose, recipient, transfer path, storage location, and retention period.

☐ The data owner approves each data item retained and records which fields were removed. The owner also records later reviews of whether the approved purpose could be met with less data.

☐ Retained evidence shows where data are protected by encryption, which encryption protocols are used, who owns the encryption keys, and how this protection covers supplier services. The record names the reviewer and gives the verification date.

☐ Data must not travel along an unencrypted route unless authorized privacy and security reviewers approve another control that provides equivalent protection and is permitted by applicable requirements. Mandatory protections cannot be waived.

☐ The access table lists every role held by a person or software account, the minimum permissions it needs, its approval, and its expiry. It includes software accounts that act without a person signing in.

☐ Testing confirms that logs capture identity, patient or context, data access, configuration, and action events.

☐ Log retention, review intervals, alerts, and thresholds have been approved and tested.

☐ The contract sets acknowledgement and notification deadlines by severity, starting from detection. It specifies what each notice must contain and who must receive it.

☐ Responsibilities are assigned for preserving evidence, investigating, providing status updates, and supporting regulatory work.

☐ The agreement specifies escalation and remedies when the supplier fails to meet an obligation.

☐ The retention record specifies how long each item must be kept, the deletion method and deadline, any legal hold that requires preservation, export requirements, and the responsible owner. It states how return or deletion duties apply to every relevant subprocessor.

☐ A certificate or technical verification provides evidence of deletion.

☐ The hospital has audit rights and a remedy for non-compliance with deletion requirements.

☐ The technical owner compares actual data movement with the approved data-flow map at the required interval and after specified changes.

☐ Every difference between actual and approved data flows has recorded evidence identifying the affected data.

☐ Each data-flow difference has a record of the immediate action to contain it and how the issue was escalated.

☐ The decision records whether the data flow will be corrected or restricted, or whether an allowable exception with equivalent protection has been approved.

☐ Verification confirms that the resulting data flow meets the authorized arrangement before the issue is closed.

G4 Vendor proof of concept

G4.1 Before the POC

☐ The sourcing record states whether a separate vendor POC is required because existing evidence or an equivalent authorized evaluation leaves important questions unanswered about suitability, performance, connection to hospital systems, or fit with the workflow. If the POC is omitted, the procurement owner and responsible governance authority have recorded the evidence and reasons, subject to local policy.

☐ Before a required POC, the evaluation plan defines the question about vendor selection or system setup that the evaluation must answer. It specifies the products and setups to test, scenarios, what results will be compared with, criteria, cost, and duration.

☐ The POC authorization defines the evaluation and its limits. It does not authorize routine operation, expansion, or clinical reliance on its outputs. Data and live-system access need separate clearance, and any live pilot requires the applicable limited operational approval and production confirmation.

☐ The POC authorization names the procurement owner, evaluation lead, evaluators, specialist reviewers, and approver. It records conflicts of interest and the start and expiry dates.

☐ Approval is in place before the POC starts for data use, access, privacy, and security controls, including the controls that keep the evaluation separate from systems or activities it is not authorized to affect.

☐ Approval is also in place before the POC starts for safety measures, incident response, rules for stopping the evaluation, and arrangements for disposing of its data.

G4.2 During the POC

☐ The log records each evaluation session's evaluator, configuration, scenario, data, and result. It includes error severity, supplier involvement, departures from the plan, approved amendments, stop events, and incidents.

☐ Evaluators remain independent. Vendor staff may support the exercise but may not select cases, change results, or decide the outcome.

☐ Where possible, the POC uses synthetic or de-identified data and connections to test systems rather than operational systems.

☐ The required separate privacy, security, safety, legal, and governance clearances are in place before PHI is used or a live system is accessed.

☐ The approved limits of the non-operational evaluation are enforced.

☐ The results identify the exact vendor, release, and configuration evaluated and document deviations, errors, and incidents.

☐ Evidence is linked to findings on workflow, subgroups, support, ability to exchange information with other systems, cost, and ending the arrangement.

☐ The results separately record differences between the evaluated configuration and the proposed production service.

☐ The evaluation record includes the result and decision for every POC criterion specified before testing.

G4.3 After the POC

☐ The procurement decision after the POC states whether procurement will stop or continue. It identifies whether more evidence or another comparison is needed and whether specifications or negotiated conditions must change.

☐ The procurement decision explains the reasons and remaining uncertainty and assigns owners and dates to any conditions. It gives no authority for operational use.

☐ The record of sourcing arrangements and responsibilities reflects the decision made after the POC.

☐ The POC's accounts, credentials, connection points, integrations, and vendor access have been removed before closure.

☐ Verification at POC closure confirms that its data and outputs have been returned, retained, or deleted as approved.

☐ Each unresolved issue affecting live use is assigned to the team handling the contract or local assurance, with a named owner and due date.

Appendix H AI incident learning record

This record covers events that caused harm and near misses that could credibly have caused harm. It links to the hospital's incident account and supporting evidence without copying them. Facts that have not been established are marked Unknown; N/A applies only when a question does not apply. A checked item means that the required facts are available or a reviewer has accepted why the item does not apply.

H1 Service identification and context

☐ The record gives the incident identifier, report date and time, reporter's role, and links to related records. It identifies the type of AI, enabled features, and exact version and setup involved.

☐ The record describes the site, service, workflow, and operating conditions at the time of the event, including tier, approved scope, and lifecycle status. It identifies the affected period and users, and any restrictions now imposed on the service.

H2 Event, impact, and detection

☐ The event account reconstructs the sequence from preserved evidence and recorded times and compares what should have happened with what occurred. It describes the inputs, data, and point reached in the workflow and identifies the first known departure from the expected process.

☐ The record identifies the types and numbers of people, records, messages, decisions, or actions affected.

☐ The event account describes the actual effects, their severity, and the time in which harm occurred or could occur.

☐ For a near miss, the event account explains the credible consequences if the event had continued or happened again.

☐ The record includes findings on care, access, rights, privacy, security, equity, workflow, employment, official records, or communications, with each finding linked to supporting evidence.

☐ The detection record identifies who or what detected the event, how it was detected, any threshold that triggered detection, and the times of detection and reporting.

H3 Containment and investigation

☐ The response record identifies the immediate steps taken to limit harm, the service's current status, review of affected records, notifications, and responsible authority. If a stop threshold was reached, it states when the affected function was suspended and the alternative workflow activated.

☐ Linked evidence shows whether the containment measures stopped or reduced the harmful behaviour. The response record identifies who or what remains at risk.

☐ The investigation has a named lead responsible for directing it, a deadline, and assigned hospital and supplier duties.

☐ The investigation has tested possible causes in the technology, data, workflow, user practices, and history of changes and recorded its findings.

☐ The findings distinguish established causes from explanations still under investigation and identify any escalation needed to obtain missing evidence.

H4 Corrective action and restart

☐ Each corrective or preventive action identifies the affected component, required release or procedure change, responsible owner, and due date.

☐ Each corrective or preventive action states how success will be tested, what result is required to pass, and who may confirm completion.

☐ Each action has a recorded date and result of the effectiveness review and any escalation required because the review was overdue.

☐ Someone independent of the corrective work verifies each action's completion before it is closed.

☐ The governance decision records the permitted scope, remaining risk, deciding authority, and date.

☐ Evidence satisfying every prerequisite has been accepted before restart.

☐ A named authority approves restart before it occurs. The decision records the first review date and any expiry.

☐ Production confirmation has been issued for the release and scope before restart.

☐ Each required patient follow-up, record correction, or action reversal has a named responsible owner, recorded completion date, and linked verification evidence.

H5 Communication and shared learning

☐ The communication record identifies each applicable notification or report, its recipient, responsible owner, due date, and completion date, with evidence of delivery.

☐ Each finding that requires changes to procurement, monitoring, staff competence, or policy is linked to a corrective action with a recorded owner, due date, and record confirming completion.

☐ Before a lesson is shared, reviewers confirm that it is de-identified and that disclosure is limited to what is necessary. The sharing record identifies who authorized sharing and when.

☐ The sharing record identifies the approved route used and links to the shared lesson, or explains why it was not shared.

☐ If another site must act, the shared record names the person taking responsibility at that site and the date they acknowledged it.

☐ Reviewers have verified that the shared lesson contains no identifiable patient information.

Appendix I Shared fields for governance records

These shared fields record the evidence needed across intake, the registry, the dossier, and later lifecycle decisions. When several checklists rely on the same report or other evidence, their records can refer to one approved, stored copy. A checked box means that the information required for the current decision is present and verified. The Hospital AI Governance Intake and Lifecycle Record provides an adaptable record for one use; the companion checklists and incident learning record provide checks against that evidence.

The record for each field or check states whether it applies, the decision for which it is required, where its evidence is stored, who verified it, and the verification date. An item is marked Met only after its evidence has been accepted for the current decision. Otherwise, it is marked Open, with a status of Pending, Failed, Unknown, or Expired. Every Open item needs an owner, due date, and a statement of the decision or activity it blocks. Evidence needed only at a later stage may remain Pending while earlier work proceeds, provided that work does not depend on it. N/A is used only when a reviewer has accepted why the item does not apply; N/A cannot remove a mandatory control. These rules apply to every field below.

I1 Use and approved configuration

I1.1 Identification and ownership

☐ The intake record identifies the AI use and the identifiers used to track it in hospital systems, the AI registry, and the workflow.

☐ The intake record identifies the supplier, product, release, and configuration, and the decision or action the system will affect.

☐ The ownership record names the clinical/end-user, operational, and internal technical owners. For each, it records authority, accepted duties, date of acceptance, and backup. It names an executive sponsor where required and identifies the governance lead and dossier coordinator separately.

☐ The approved-use record specifies who may use the system, at which sites, in which settings and shifts, for which populations, and with which exclusions. It records the expected volume and maximum approved scale.

☐ The intended-use statement describes who uses the system and why, its inputs, the output or action it produces, and who receives it. It records the level of AI involvement for each task: Abstain, Suggest, Draft, Act with direct oversight, or Carry out defined tasks independently. It identifies the affected workflow, human decision authority, and the status of any resulting record, such as draft or final.

☐ The approved-use record lists prohibited uses and uses outside the approved scope separately.

☐ The dossier index records the registry link, dossier location and version, local form identifier, and dossier coordinator. The governance lead directs records to the required reviewers and maintains the status of the use. Maintaining these records does not itself give anyone authority to approve the use.

☐ The dossier index links the approved user instructions, data map, assurance record, and hazard record.

☐ The dossier index links the monitoring, competency, and transparency records.

☐ The dossier index links the contract, change, incident, and retirement records.

I1.2 Tiering and initial screening

☐ The tiering record shows the assessment from Tier 4 down to Tier 1 and explains the most serious credible consequence of the defined use. It describes how a failure could cause harm, how quickly harm could occur, whether it could be detected or reversed, and how widely it could affect people or services.

☐ The scope record identifies the limits and required human decisions that define the use. It links evidence that these restrictions are in place, enforced, and verified.

☐ The tiering record states the proposed and confirmed Tier 1–4 classifications separately, with evidence, uncertainty, reviewer, decision authority, date, and reassessment triggers. Planned or untested safeguards have not been used to lower the tier. The risk remaining after controls is documented separately.

☐ The Tier 0 record states the next activity requested and records evidence for every criterion, the reviewer, and date. It lists every failed or unresolved item.

☐ The Tier 0 record states the decision—Reject, Hold, or Proceed—and the exact activity it permits. For each condition, it identifies the owner, due date, required evidence, and who may verify completion. Prohibited uses or unacceptable features that cannot be corrected receive a Reject decision. Any activity that depends on missing essentials or an unmet prerequisite remains on Hold.

I1.3 Data and configuration

☐ The data inventory records each input's source, sensitivity, whether it includes patient data, and the form in which those data are held.

☐ The data inventory records each input's purpose, the authority permitting its use, and how its use has been approved.

☐ The data inventory names the owner responsible for each input's quality or lineage. Lineage is the record of where the input came from and how it was changed. The inventory documents each input's retention rule and known limitations.

☐ The data inventory confirms that each input meets the minimum-necessary rule: it contains only data needed for its approved purpose.

☐ The data-flow map identifies every data source, element, recipient, transfer, and connection point.

☐ Authorized testing compares the proposed data-flow map with observed data movement before live access is granted or the service is released. The findings and approved map are retained.

☐ The data-flow record identifies every storage and log location, hosting region, supplier, and subprocessor. A subprocessor is an organization engaged by a service provider to process data on its behalf as part of the hospital service. Each organization's legal role under the applicable law is confirmed.

☐ The data-flow record states whether personal health information (PHI) leaves the hospital. It records retention and deletion arrangements, any permitted use beyond the original purpose, and the authority for access. It includes the date and result of the latest comparison with actual data flows and any unresolved difference.

☐ The integration record describes how systems connect and the agreed structure and format of the data they exchange. It records the user or software identities and their permissions, the person responsible for release, possible failures, and how they will be monitored.

☐ The continuity plan describes the alternative workflow and links test evidence of the volume it can handle and how long it can operate.

☐ The continuity plan explains how outstanding work and inconsistent records will be identified and resolved. It describes how to restore an earlier approved setup or safely disable the service and activate an approved alternative. It records the recovery target and the result of the latest exercise.

☐ The interface specification defines the structure and format of inputs, outputs, or actions and where they are sent. It states where resulting records are placed and their status, such as draft or final. It describes how users will see the information source and any uncertainty.

☐ The interface specification states when user approval is required, which later uses of the output are prohibited, which fields are kept for audit, and what happens if an output is wrong or missing.

☐ The configuration record identifies the versions of the model, prompts, templates, information sources used for retrieval, software components, system connections, policies, and supplier releases that apply to the use. It identifies their configuration group. Together these describe the configuration baseline: the setup being proposed or approved. The proposed and approved setups are clearly distinguishable.

☐ The configuration baseline names who may authorize changes and states the notice periods, routine and emergency change procedures, and the rollback target or approved plan for safely disabling the service.

☐ The configuration record includes the date and result of the latest check that the actual setup matches the approved baseline.

☐ The configuration-group record identifies the defined use, the variants treated as equivalent, and the supporting evidence. It states the approved operating limits for population, context, interfaces, data, permissions, versions, and workload. It records both the approved maximum number of groups for the service and the number currently supported.

☐ For each configuration group and variant within it, retained evidence shows that the hospital has capacity for validation, staff competence, monitoring, support, and rollback or safe disabling within the approved operating limits.

☐ The configuration-group record gives the reason for the limit on configuration groups and names who approved it. For an exception, it records funding, approving authority, and expiry. It states whether reaching the limit requires groups to be combined, an existing group to be retired, or support capacity to be increased. An exception cannot remove core controls.

Different setups for the same defined use may form one configuration group only when evidence shows that they can be evaluated, monitored, and supported together. Evidence must also support sharing training, assurance, and arrangements for rollback or safe disabling across those setups. A common product name, site, or version is not enough. The proposed default of one group per service is a planning limit, not a scientifically established safety threshold. A hospital choosing a different limit must show that it has the capacity to support it.

☐ Tests with intended users of the patient-facing service show whether they understand the information and can use the required route to obtain further help. The test record includes accessibility findings and states whether older adults participated.

☐ The patient-facing service record describes arrangements for caregivers or people authorized to act for the patient, alternatives involving a person or a non-digital route, and how staff respond to outside AI advice brought by patients. An appropriate alternative is available where feasible and whenever required by applicable duties or approval conditions. If an alternative is unavailable, the record explains why and how concerns will be addressed; an unavailable mandatory alternative prevents the affected use from proceeding. The record links the responsible owners, monitoring measures, response thresholds, and the patient, equity, accessibility, and Indigenous governance review in the intake record.

I1.4 Value, hazards, and human oversight

☐ The benefit record defines how each expected benefit will be measured, the data source, the population or events counted in the denominator (the total used to calculate a rate), and the baseline or comparison period.

☐ The benefit record specifies the target, target date, and required subgroup reporting.

☐ The benefit record names the owner and sets the review interval.

☐ The value assessment reports the benefit achieved to date and the full cost of introducing, operating, and retiring the service. It states what action is required if the expected benefit is not achieved.

☐ The hazard record describes, for each hazard, the output or action that could fail or be misused, how this could cause harm, who could be affected, how quickly harm could occur, and whether the problem could be detected.

☐ The hazard record documents the measures intended to prevent or detect each hazard, their implementation status, and links to verification evidence.

☐ Warning, action, and stop criteria are set for each hazard, with response owners, deadlines, and restrictions. Meeting a stop criterion requires immediate suspension of the affected function, activation of fallback, and incident escalation.

☐ The hazard record states the decision about the risk that remains after controls, the remaining uncertainty, and who made that decision.

☐ Where human review is required, the review record names the reviewer and links evidence of competency. It specifies the source information shown, time allowed, and assumed workload.

☐ The review record states the reviewer's authority to reject an output or proposed action, how they must indicate approval, and what audit evidence will be kept. It specifies how to raise a disagreement, the alternative workflow, and the response time. Any action allowed without prior human confirmation has explicit, separately approved limits on what the system may do.

☐ The transparency record states whether requirements to inform people about the AI use apply and who has authority to decide this. Where transparency requirements apply, it specifies the approved audience, information, timing, format, language, and delivery method. It records the practical alternative offered where feasible or required, any limits, and how concerns will be addressed.

☐ Routes for questions, correction, and human review are available. The transparency record names the responsible owner and records the test result, effective date, and expiry of these arrangements.

I1.5 Status and funding

☐ The lifecycle record states one current status: Research only, Intake, Local assurance, Approved but not live, Early-adopter implementation, Routine operation, Restricted, Paused, or Retired. It separately records the status start date, the relevant decision stage, and the decision requested. Expansion belongs to review of continued use and changes to approval: it is a decision about scope, not a separate status. The six recurring governance tasks describe the work supporting these decisions.

☐ The function inventory identifies the AI type and every enabled function: capturing information, retrieving it, creating content, making recommendations, writing records, communicating, or taking an external action.

☐ Each component is classified as build, buy, or partner, with any mixed arrangement explained component by component.

☐ The early-adopter plan identifies the initial group, maximum starting scope, required user practices, support arrangements, and first review. It states which aspects of the service are proposed for later expansion.

☐ The funding record states funding for one-time implementation separately from recurring operation, with the source and duration of each commitment.

☐ The funding record specifies the committed amount or capacity, accountable owner, and renewal-decision date.

☐ The funding record confirms coverage for licensing or hosting, monitoring, training, support, incident response, continuity, and exit. It states what happens when funding ends.

I2 Procurement and approval records

I2.1 Procurement choices and evidence

☐ For each material component, the sourcing record states the criteria and evidence supporting the choice to build, buy, or partner, with the reviewer and date. It records the scores if the selected comparison method uses numerical scoring.

☐ Before commitment, the sourcing record compares the resources and cost needed throughout the service's life, its dependencies, and whether the arrangement can be ended safely. It records exceptions and the reasons for the final choice.

☐ The sourcing record separately states whether a vendor proof of concept (POC) is required. If it is omitted, the evidence and reasons accepted by the procurement owner and responsible governance authority are retained, subject to local policy. An internal build without vendor selection may mark the vendor POC N/A; local assurance remains required.

☐ Where a vendor POC is required, its record states the purchasing question it must answer, the exact setup tested, limits on data use, evaluation criteria, duration, and expiry. The results are added when it ends.

☐ The POC closure record states the procurement decision and confirms that the evaluation itself granted no operational authority.

☐ The POC authorization identifies the responsible roles and decision-makers, reviews required before it starts, evaluation procedure and criteria, and rules for stopping it.

☐ The POC log records evaluation sessions, amendments, vendor involvement, stop events, and incidents.

☐ The post-POC record confirms the results and procurement decision and names the independent verifier.

☐ The post-POC record confirms that POC access has ended and its data have been returned, retained, or deleted under the approved arrangements.

☐ The post-POC record identifies the responsible owner and date for each handover of work or records.

☐ The evidence inventory identifies each source by version and date. It states how independent the source is of the supplier or team proposing the use and records declared conflicts of interest.

☐ The evidence inventory states whether each item applies to the release, population, setting, workflow, and output under review.

☐ The evidence inventory explains what remains uncertain for each evidence gap. It records the reviewer, assessment decision, person responsible for obtaining or addressing the evidence, due date, and decision that must wait.

☐ Accepted evidence is linked for every evidence gap that has been closed.

☐ The assurance record identifies the approved data-flow map version and the result of its latest comparison with observed flows.

☐ The assurance record states the privacy and cybersecurity screening rule, reviewer and date, whether further assessment is required, and the reasons for the decision.

☐ The assurance record states each required assessment's findings, restrictions, and decision.

☐ The assurance record states each condition's owner, due date, the decision or activity it affects, and who may verify completion. It links the evidence accepted to close the condition when that evidence becomes available.

☐ The contract record specifies intended and prohibited uses, rights to data, access to evidence and logs, and update-notice requirements.

☐ The contract record states the required monitoring and incident-response service levels, duties for correcting problems, liability, insurance, continuity, audit arrangements, and arrangements for ending the service.

☐ The contract record also states retention and deletion duties, hospital decisions, exceptions, sign-off, and expiry.

I2.2 Validation and approval

☐ The validation record identifies the production release and configuration, data path, population, and evaluation period tested.

☐ The validation record describes the cases included, the comparator or reference standard used to judge results, the workflow, and the intended users.

☐ Validation criteria are set before testing for each applicable area: benefit, performance, serious errors, subgroups, and equity. The validation plan specifies how uncertainty is handled and the rules for Pass, Fail, or Insufficient evidence. Critical criteria are explicitly identified, and every applicable critical criterion must Pass. A criterion is marked N/A only when a reviewer accepts that it does not apply.

☐ The validation report presents results and deviations and records the reviewer, date, and decision.

☐ Tests have assessed how representative users work with the system across the relevant sites and shifts. They cover normal and difficult work, time pressure, disagreement with the system, errors, and use of the alternative workflow.

☐ Acceptable results for task completion, time taken, errors, workload, and accessibility are set before testing how people work with the system.

☐ The human-factors report documents results, workarounds, automation bias (relying on the system when it should be questioned), and corrective actions. It records the reviewer, date, and decision to pass or hold.

☐ For each relevant role, the competency record states what staff must learn, how they will learn and be assessed, the pass criteria, assessor, and result. It states when competency evidence expires or what triggers refresher training or reassessment and what access is restricted while the requirement is unmet.

☐ The rollback or manual-continuity test record states the trigger, authority, recovery target, result, and unresolved defects.

☐ Governance approval is recorded with a named authority and date, stating what preparation may proceed and its conditions.

☐ Production confirmation is recorded as a separate decision with a named authority and date. It authorizes release only after the approved version, scope, controls, and every live-use prerequisite have been verified.

☐ The approval record states the approved baseline, permitted and prohibited uses, risk remaining after controls, and remaining uncertainty.

☐ Every decision records its conditions, monitoring requirements, expiry, next review, and escalation route.

☐ Each condition has an owner, due date, the decision or activity it prevents, and an authority to verify completion.

I2.3 Configuration by AI capability

☐ For predictive or other non-generative systems, the configuration record identifies the outcome, category, or task to be assessed, the reference standard for judging results, data source, version of the calculation, comparison baseline, and decision threshold. The Monitoring operations fields record ongoing monitoring.

☐ For non-generative systems, the applicable acceptance criteria for task performance, calibration, serious errors, and subgroup results are set. Operating rules for warning, action, stopping, and restarting are recorded separately from validation results.

☐ For generative systems, the configuration record specifies the approved and prohibited data, information and circumstances the system may use as context, outputs, and actions. It identifies the exact setup to which these rules apply.

☐ For generative systems, the configuration record defines the types of error, limits on severe errors, and required checks against source information. It records the status of outputs, such as draft or final, how data move, how long they are kept, and which cases will be sampled for monitoring. A control is omitted only when a reviewer has accepted why it does not apply.

☐ The baseline names who may respond to a breached threshold, disable or roll back the service, and approve configuration changes.

☐ For action-taking systems, the configuration record states the goal, permitted stages of work, and rules for moving between them (the state machine). It identifies approved tools and actions, user or software identities and permissions, and how each activity is linked to the correct patient or context.

☐ For action-taking systems, the configuration record states when human confirmation is required and any approved permission to act without it. It defines the memory rules, including those for information retained between interactions, limits on actions and repeated attempts, protection of audit records, handling of exceptions, and how work can continue manually.

☐ The baseline names who may prevent the system from changing data or other systems and explains how partly completed actions will be resolved.

☐ The agent test record links each deliberately challenging input or attempt to defeat safeguards to its pass criteria and results.

☐ The agentic monitoring record defines thresholds specific to each applicable action and names the people responsible for responding and their deadlines.

I2.4 Verification before live use

☐ Before live use, the exact proposed version, connections to other systems, and data flows have been tested against acceptance rules set in advance. Testing uses the intended operating environment or an authorized equivalent. Any material difference from the live environment is assessed before release.

☐ Tests cover the serious errors specified in advance and assess the required subgroup performance within the intended workflow.

☐ Test results show how capacity, monitoring, the alternative workflow, and either rollback or safe disabling perform against the approved criteria. The results are retained.

☐ Verified results show that incident-response and recovery exercises meet their approved criteria.

☐ The verification record lists defects, responsible owners, and dates. It states any conditions for releasing the service in stages and identifies who may accept the results.

☐ Every applicable critical criterion and mandatory live-use condition is met before production confirmation.

I2.5 Responsibility and sign-off

☐ The responsibility table identifies who controls each component, who performs the work, and which hospital owner is accountable for each decision the hospital retains.

☐ The component responsibility matrix assigns responsibility for changes, evidence, monitoring, investigation, and correction. It records the applicable service levels, including response or availability commitments, together with funding and duties for rollback or safe disabling, continuity, and exit.

☐ The decision record states separately what the evidence supports, the approved baseline, governance approval, production confirmation, and any approval to expand the use. Initial release, material release changes, expansion, and restart each require production confirmation for the scope concerned.

☐ The scribe record states the approved setup and every disabled feature that could use its output or take further action.

☐ Before sign-off for an action-taking system, verification confirms that the controls specified in its baseline are in place and working. These cover permitted actions and authority, access permissions, retained information, audit records, and recovery.

☐ Specialist conclusions and sign-offs identify the exact product, version, use, and decision they cover. Each required finding is linked to the activity that depends on it.

☐ Every live-use condition is resolved before production confirmation.

☐ The acknowledgement record states each owner's name, role, authority, decision, and date.

☐ Final sign-off is linked to evidence that the committee had the required attendance to decide, or to the decision made under an authorized delegation.

☐ The decision record names the person authorized to verify and close each condition.

☐ For a transfer from research to operational use, the transfer record states the research ethics reference, proposed operating scope, evidence that can be reused, and remaining gaps. It records who may authorize operation, the support and supplier duties, and the accountable owners' acceptance.

I3 Monitoring, change, and retirement records

I3.1 Monitoring operations

☐ The monitoring record names the accountable owner and backup, report location, and data cutoff—the latest date and time included in the report. It sets targets for data completeness and reporting delay.

☐ Results from testing the full monitoring process and the fallback during an outage are recorded.

☐ For each monitoring measure, the record states what is measured, the source, denominator, calculation version, baseline, and comparison period. It states the minimum data needed, how uncertainty will be handled, and any required reporting by site or subgroup. The response to missing or late data is decided in advance.

☐ Each measure's record names its reviewer and backup, how often they review and report it, and who receives the report. It records the last and next review dates and what happens if review is overdue. The current result is linked to the record specifying the response at each threshold.

☐ The monitoring record links the hazard, preventive safeguard, monitoring measure, and known detection gaps for each applicable failure pattern: target and representation bias; data and label noise; covariate shift or out-of-distribution inputs; prevalence changes or label shift; concept drift or target-definition changes; data leakage or shortcut learning; prediction or calibration errors; data-pipeline or integration failures; hallucination, generation, or retrieval errors; agent goal or execution failures; reward hacking or specification gaming; security or privacy failures; automation bias, human, or workflow failures; and monitoring or response failures. It records the approved comparator against which results will be assessed, the independence of the evidence sources, and any accepted exclusion.

☐ The sampling record states random and targeted sample selection separately. It includes relevant cases that did not trigger an alert, as well as excluded, failed, and abandoned cases. It specifies sample size or selection rule, subgroup coverage, review frequency, reviewer, and response to an unexpected finding.

☐ The monitoring record describes how the labels or other results used to judge the system are checked and when actual outcomes become available. It explains how missing data, disputed labels, changing definitions, or too few observations affect interpretation and the decision to continue use.

☐ The monitoring record links the routes for patients and staff to report concerns, checks that required safety conditions are met, deadlines for assessing and escalating concerns, and the latest authorized exercise testing the response to monitoring findings. It records resulting changes to the hazard list, monitoring plan, and tests.

I3.2 Measures and responses

☐ Performance, use, and drift are measured and reported separately. Performance describes the system's results; use describes how it is used; drift describes departures from the conditions supporting approval, including changes in data, population, workflow, or system behaviour.

☐ Records for performance, use, drift, outcomes, and equity measures include all Monitoring operations fields. They link to shared definitions and results so these need not be copied.

☐ Every measure's record states the target, applicable warning, action, and stop thresholds, current result, responsible owner, and required response deadline.

☐ The outcome and equity assessment states whether subgroup analysis applies, which local groups are relevant, and where material data gaps remain.

☐ Outcome and equity measures use the same monitoring fields. Their reports include sample or event counts, uncertainty, and missing data for every required group.

☐ The outcome and equity assessment compares current results with the targets or minimum acceptable results for each subgroup. It names the reviewer and person responsible for mitigation and records the review date and required response.

☐ The response plan defines warning, action, stop, and restart criteria and explains how uncertainty, insufficient data, or reporting delay affect the decision. It distinguishes criteria that trigger action during operation from criteria used to accept the system in validation.

☐ The response plan specifies the required action, authorized primary and backup responders, deadline, and escalation route for each trigger. Reaching a stop threshold requires immediate suspension of the affected function, activation of fallback, and incident escalation.

☐ The response plan specifies fallback, communications, and any required review of affected records.

☐ Any temporary arrangement to maintain service during an emergency has separate authorization, a time limit, and equivalent protection. Mandatory duties still apply.

☐ The response plan specifies the evidence needed to continue under an approved restriction or to restart. Restart requires accepted evidence that the required work is complete, a decision authorizing the scope of use, and production confirmation for the exact release and scope.

I3.3 Incident readiness and reporting

☐ The incident record explains how to report an incident, which services and hours are covered, how severity is assigned, and who provides on-call backup.

☐ The incident arrangements set deadlines for acknowledgement, containment, investigation, notification, and correction.

☐ The incident arrangements assign responsibility for preserving evidence and for supplier investigation.

☐ The incident plan names who may suspend or restart the service and the requirements for each decision. It retains the mandatory response when a stop threshold is reached: immediate suspension of the affected function, activation of fallback, and incident escalation. It requires production confirmation before restart.

☐ The readiness record states the results of the latest exercise, defects, and escalation of overdue actions.

☐ The governance reporting arrangements identify recipients, the maximum reporting interval, data cutoff, preparer, and reviewer.

☐ Governance reports show each use's tier and status, departures from thresholds, incidents, and material changes.

☐ Reports covering all hospital AI uses also identify heavy dependence on particular suppliers, benefits achieved, equity findings, reviews waiting to be completed, and overdue conditions.

☐ The reports identify the decisions needed and record each action's owner, date, and completion status.

☐ Incident follow-up assigns an owner and deadline to every corrective and preventive action and specifies how its effectiveness will be checked.

☐ The incident follow-up record documents required corrections to affected records and the restart decision.

☐ The incident follow-up record states whether lessons have been documented, stripped of identifying information, and shared through the approved route.

I3.4 Change and record updates

☐ The change record shows the currently approved and proposed baselines together so reviewers can compare them.

☐ The change record identifies every component, workflow, population, permitted-action limit, and supplier arrangement affected by the change.

☐ The change record explains why the change is proposed. It states the impact and tier assessments, reviews that must reopen, required evidence, and acceptance criteria.

☐ The change record states the decision authority and date, arrangements for rollback or safe disabling, and monitoring after the change.

☐ The approval table identifies the specialist reviewers and decision-makers required for the tier, enabled functions, and requested decision. It includes governance and production authorities, those who may authorize exceptions, and any external authorities. It states when each prerequisite must be met.

☐ Every decision record states who approved the decision, the source of their authority, any required committee attendance or delegation, the date, expiry, and any disagreement with the decision.

☐ The decision record names the person authorized to verify and close each condition.

☐ The affected decision or activity remains on hold while a mandatory approval is missing. Unrelated preparation can continue only if the recorded scope explicitly permits it.

☐ After every lifecycle decision, the update record states who updated the registry and dossier, when they did so, the decision link, and current status. This applies to governance approval, production confirmation, expansion, material change, restart, restriction, pause, and retirement.

☐ Each update records the approved baseline, tier, owners, conditions, monitoring plan, and next review.

☐ Each update links the relevant supplier, contract, change, incident, restart, and retirement records and identifies any version it replaces.

I3.5 Retirement and learning

☐ The retirement record explains what triggered retirement and why it was chosen. It records the deciding authority and date, continuity plan, and communications.

☐ Retained evidence confirms that access, credentials, integrations, and scheduled work have been disabled.

☐ The retirement record confirms that outstanding actions and affected records have been checked and resolved, or transferred to the replacement workflow.

☐ The retirement record verifies completion of contractual obligations and requirements for returning, retaining, or deleting data.

☐ At retirement closure, the registry status is updated and lessons for other uses are recorded.

☐ Every unresolved action at retirement closure has an owner and due date.

Appendix J Evidence sources and local adaptation

J1 Purpose and limits of the evidence

This framework draws on standards, official guidance, peer-reviewed research, published accounts of implementation, and consultation with health-AI leaders. These sources informed the proposed hospital process: who may make a decision, what evidence they need, and how the decision and its conditions are recorded. This appendix explains what each source contributed, where it applies, and what authority it carries. It also explains how the sources were combined into the proposed process.

The initial review of source versions and status was completed on 16 August 2026, followed by targeted updates on 28 September. Publicly available reference details and cited claims were checked again on 1 October 2026. Hospitals must check guidance that may have changed before relying on it locally. The framework is an operational proposal, not a systematic review or legal opinion. It cannot replace regulatory authorization or a medical-device quality management system, which governs the processes and records required to develop and maintain the device.

For each source used in a local decision, the dossier should state which claim it supports, where it applies, and what authority it carries. It should also identify conclusions that still require local evidence. Binding law and approved hospital policy create obligations within their scope. Voluntary standards and implementation studies usually offer methods, principles, or experience that need to be interpreted for the proposed use. Reviewers should distinguish what the source establishes from what remains unresolved. Meeting one law, standard, or policy does not establish compliance with another or show that the service is suitable for the hospital's intended use.

Must marks either a core requirement of this proposed governance model or a binding obligation confirmed as applicable in the hospital-specific appendix. Should marks a recommended approach; a hospital may use an alternative that meets the same objective. May marks a permitted option. These words do not give guidance legal force outside its jurisdiction. The responsible legal, regulatory, professional, or policy owner must determine which obligations apply.

Hospitals need to justify their choices about committee structure, risk scores, review intervals, and monitoring thresholds. The rationale needs to reflect the proposed use, plausible consequences, available evidence, local capacity, and applicable law. The literature does not establish one approach that suits every hospital. Regulatory authorization, external validation, certification, and supplier quality systems can support a decision when they cover the proposed service. If the population, workflow, configuration, or purpose differs, reviewers must determine whether that evidence still applies. Local review cannot replace authorization required by law.

Human review must be tested in the workflow where it will occur. The reviewer needs relevant knowledge, enough information and time, authority to disagree, and a practical alternative course of action. A professional title or approval button cannot establish those conditions, particularly under time pressure or when an output is difficult to verify. Longer-term effects on skills also need attention. Experienced staff may lose skills, while learners may fail to acquire them when AI mediates their training. Evidence on these effects is still developing, so hospitals should preserve and observe the skills needed for safe unaided work without assuming that harm will occur or that competence will remain unchanged.

The sources support the framework's design choices, but the combined governance process has not yet been evaluated as an intervention. Hospitals adopting it should assess how well the process works in practice. This includes checking whether the inventory is complete, decisions are timely and supported by evidence, and approval conditions are completed. They should also examine how problems are detected and handled, whether

retirement is verified, and how much work governance requires. Hospitals should also measure effects on safety, equity, service performance, and patient outcomes and use the findings to guide revisions.

The source review considered authority, relevance, and recency, checked versions, dates, and status, and examined whether each claim was supported in the source's setting. Local teams should repeat checks where guidance may have changed and retain enough detail in the source register to trace earlier interpretations.

J2 Contribution of this framework

Earlier work provides principles, methods, and examples of governance in practice. This framework connects them to hospital decisions and the records needed to support those decisions. The International Organization for Standardization (ISO), European Union (EU), and National Health Service (NHS) sources serve different purposes and have different scopes. **Table J1** explains what each source contributes, which implementation questions remain for hospitals, and how this framework addresses them. The notes after the table explain the assessments that hospitals still need to make locally.

The NIST AI Risk Management Framework [42] provided concepts of governance, context, measurement, and risk response that we adapted into hospital records and decision responsibilities. Its Generative AI Profile [12] informed the review of general-purpose services, third-party suppliers, and connectors. The 2026 report NIST AI 800-4 [91] provides terminology for monitoring deployed systems and describes limitations in validated methods and common practice.

ISO/IEC 42001 [47], a joint standard of ISO and the International Electrotechnical Commission (IEC), contributed management-system approaches to responsibility, documentation, evaluation, and improvement. ISO/IEC 42005 [92] informed impact assessment. The framework adapts these ideas without reproducing the standards or establishing ISO certification.

World Health Organization (WHO) guidance on AI for health [6] and large multimodal models [11] informed the requirements concerning human rights, patient interests, transparency, and institutional responsibility. The FUTURE-AI consensus guideline [93] addresses trustworthiness throughout the healthcare AI lifecycle. The International Medical Device Regulators Forum (IMDRF) framework for software as a medical device (SaMD) [44] and NICE Evidence Standards Framework [10] helped relate evidence expectations to what a system does and its credible consequences. The Tier 0 screen and Tier 1–4 classifications were developed for this hospital governance framework; they are not device classes, statutory high-risk categories, or external certifications.

Published health-system programs offer operational precedents for this process. Nong and colleagues [94] documented variation across academic medical centres. Duke's Algorithm-Based Clinical Decision Support (ABCDS) program [15] uses registration, risk-based review, and checkpoints between lifecycle stages; Wisconsin's Clinical AI and Predictive Analytics (CAIPA) Committee [95] connects intake and multidisciplinary review with monitoring and retirement. Stanford's Fair, Useful and Reliable AI Models (FURM) assessment [96] considers ethical fit, simulated usefulness, financial sustainability, implementation feasibility, and prospective monitoring. The Canadian health-system case [97] informed local evaluation and implementation. Kim and colleagues' people, process, technology, and operations (PPTO) framework [98], published in February 2026, informed consideration of the organizational capabilities needed to carry out this work.

The Coalition for Health AI (CHAI) Governance Playbooks [82] informed oversight and monitoring across all the hospital's AI uses. Whittaker and colleagues [88] described how Māori and consumer participation shaped governance in their setting. NHS DCB0160 [99] and ambient-scribing guidance [51] informed systematic identification and management of clinical hazards, as well as clinical documentation practice. A NEJM AI

editorial [13] also identified general-purpose and unregistered or unapproved AI uses, often called shadow AI, that a review centred on products may overlook.

We combined these approaches into a process for clinical and administrative hospital uses, including workforce, education, and general-purpose AI. The National Academies' Best Care at Lower Cost [1] and Friedman and colleagues' Toward a science of learning systems [2] establish the basis for learning from care. They do not prescribe the institutional model proposed here.

The framework's contribution is a common set of hospital decisions and companion records that connects these established approaches. Tier 0 asks whether a specified activity may proceed. The four consequence tiers determine how much review is needed; a separate assessment records the risk remaining after safeguards. Governance approval defines the permitted use and its conditions. Production confirmation then checks the actual release and workflow before live use. One dossier names the owners and the evidence due at each decision, following the use from procurement and evaluation through monitoring, change, incidents, and restart to expansion or retirement. Within this common process, function-specific checks address predictive AI, generative AI, scribes, and agents.

The minimum and comprehensive models, shared-network arrangements, and hospital-specific appendix help hospitals assign these functions according to their resources and authority. Expansion decisions explicitly consider support capacity, funding, and the number of configurations that need separate oversight. Consultation across all 15 TAHSN institutions and with other international health-AI leaders identified operational needs and helped refine sign-off, literacy requirements, retirement checks, and the illustrative cases.

The HALO Model, published in June 2026, offers a related conceptual account connecting an adapted AI lifecycle with learning-health-system cycles and continuing assurance. [100] Both approaches emphasize learning from operation and revisiting safeguards as conditions change. In this framework, the figures and companion record show who does that work, when a decision is needed, and what evidence supports it.

Table J1 compares each source's contribution and scope with the hospital implementation questions it leaves unresolved. The final column identifies the procedures and records this framework adds to address those questions.

***Table J1**. Contributions of prior work, remaining implementation gaps, and what this framework adds.*

Prior work	What it contributes	Questions left for local implementation	What this framework adds
NIST AI RMF 1.0 [42], AI RMF hub [101], Playbook [102], Generative AI Profile [12], and AI 800-4 [91]	Govern, Map, Measure, and Manage functions; suggested governance actions; generative-AI risks; and categories and limitations of monitoring.	These resources are deliberately adaptable across sectors. They leave hospitals to define who can authorize a use, what evidence is sufficient, and how monitoring findings lead to action within a local decision process.	Names the hospital decision-makers and keeps evidence, approval conditions, and follow-up actions in one dossier. Monitoring plans state when findings require reassessment or incident response. The record continues through controlled change and retirement (**Appendices B, E, H, I, and L**).

Prior work	What it contributes	Questions left for local implementation	What this framework adds
ISO/IEC 42001:2023 [47] and ISO/IEC 42005:2025 [92]	AI management-system responsibilities and continual improvement; assessment of effects on people, groups, and society.	Management and impact-assessment standards do not by themselves define what a hospital may use, which evidence is due at each decision, or what happens when a condition is unmet. Hospitals must specify these for the service and its clinical or operational setting.	Records impact findings and named responsibilities in the approval record for each use. Every condition identifies an owner, evidence of completion, and the consequence of non-completion. The hospital-specific appendix retains applicable standards obligations (**Appendices A, I, J, and L**).
WHO guidance on AI for health [6] and large multimodal models [11]	Ethical principles concerning autonomy, safety, transparency, accountability, inclusion, equity, and sustainability.	WHO gives ethical guidance and recommendations to several actors, including healthcare providers. For a particular use, hospitals must still determine whose interests need representation, what oversight must demonstrate, and which findings would prevent or limit deployment.	Requires evidence of how participation and patient communication shape the use, how subgroup effects are assessed, and whether human oversight works in the intended workflow. The dossier records the resulting approval conditions and matters needing continued review (**Appendices D, F, I, and L**).
FUTURE-AI international consensus guideline [93]	Lifecycle recommendations for trustworthy and deployable AI, including local validation, risk management, documentation, training, usability, and monitoring.	The recommendations cover trustworthy development and use of healthcare AI tools, including governance and monitoring. Hospitals still need to apply them through a common approval and release process across all their AI uses, specifying the evidence needed for each decision.	Sets out a hospital process with named reviewers and evidence requirements at each decision. Governance approval defines the permitted use and conditions; production confirmation checks the actual release and workflow before live use (**Appendices B, D–F, I, and L**).

Prior work	What it contributes	Questions left for local implementation	What this framework adds
IMDRF SaMD risk-categorization framework [44] and Good Machine Learning Practice principles [16]	Links evidence expectations to a system's role in healthcare decisions and the seriousness of the situation; emphasizes representative data, human–AI performance, and lifecycle oversight.	These sources concern medical devices and already address lifecycle and human–AI performance. Their scope does not cover the hospital's full range of administrative, workforce, educational, and general-purpose uses, or determine the local legal status of each proposed service.	Applies consequence assessment to the full hospital AI inventory while keeping the medical-device determination separate in the hospital-specific appendix. Reviewers assess evidence for the intended population and workflow and record the authorization needed for the proposed use (**Appendices C, D, I, J, and L**).
Health Canada guidance for machine-learning-enabled medical devices [103]	Evidence expectations for new or amended Class II–IV device licence applications, including the tested system and a predetermined change-control plan where applicable.	The guidance addresses manufacturer evidence and lifecycle controls for medical devices. Hospitals must still establish whether that evidence covers their actual population, workflow, and configuration, and who will authorize local use and respond to supplier changes.	Brings applicable device requirements and manufacturer evidence into the hospital's intended-use and local-assurance review. The dossier identifies the tested product and version, change evidence, surveillance responsibilities, and corrective actions. The hospital must also confirm that any required regulatory authorization covers the proposed use (**Appendices I, J, and L**).
Canada's National Artificial Intelligence Strategy: AI for All [23] (4 June 2026)	Canadian policy priorities for responsible adoption, skills, sovereignty, and infrastructure, with health and life sciences identified as a priority.	A national strategy sets adoption and investment priorities; it does not determine whether a hospital service is ready for use. Hospitals must work out what those priorities require for staffing, data stewardship, procurement, and continuing support.	Uses the hospital-specific appendix to record relevant initiatives and what they mean for education, data stewardship, procurement, and sustainable support. This guides implementation choices. Each proposed service still needs documented evidence, resources, and conditions for authorization (**Appendix J**).

Prior work	What it contributes	Questions left for local implementation	What this framework adds
CHAI Governance Playbooks [82] (27 May 2026)	Adaptable organizational guidance for multidisciplinary governance, data and third-party controls, impact assessment, lifecycle management, education, transparency, and feedback.	The adaptable playbooks cover organizational resources and lifecycle governance. Each hospital must still decide how to assign those functions, which evidence and signatures are required for each use, and how support capacity limits deployment or expansion.	Offers minimum, comprehensive, and shared-network options for assigning the functions, while assessing support capacity separately from consequence tier. One dossier names reviewers, records evidence and approval conditions, and tracks their completion. Production confirmation checks the actual release and workflow (**Appendices B, C, I, and L**).
Singapore MOH and HSA AIHGle 2.0 [8] (10 March 2026)	Developer, deployer, and user responsibilities; registry, local testing, deployment, monitoring, incidents, exit, and communication. Includes generative and direct-to-consumer AI and a deployer toolkit.	AIHGle provides detailed lifecycle responsibilities and a deployer toolkit for Singapore's clinical and clinical-operations settings. Hospitals need to adapt it to their local authority. They also need consistent approval and production decisions across their wider range of AI uses.	Uses one dossier for clinical, administrative, workforce, education, and general-purpose AI. It combines the registry entry, function-specific checks, patient communication, and lifecycle responsibilities with records of the authorized use, evidence, accountable reviewers, and production confirmation of the actual release and workflow (**Appendices B, F, G, I, and L**).
NHS DCB0160 [99] and NHS England ambient-scribing guidance, version 3 [51]	Clinical risk management for health IT and scribe-specific practices covering safety leadership, hazards, privacy, training, output review, incidents, and continuity.	DCB0160 addresses clinical risk in health IT; the scribing guidance addresses a particular application. Hospitals must integrate these safety records and controls with their wider AI approval process, and determine which NHS requirements apply in their jurisdiction.	Brings hazard assessments, named safety reviewers, workflow tests, fallback, and incident response into the hospital decision record. Scribe-specific production checks verify the actual service before live use. Monitoring and change decisions keep the associated safety evidence available for continued review (**Appendices D, F, H, I, and L**).

Prior work	What it contributes	Questions left for local implementation	What this framework adds
EU AI Act [43] and Commission FAQ [106]; FDA human-factors [107] and clinical decision-support guidance [108]; NICE Evidence Standards Framework [10]	EU legal requirements; U.S. guidance on human factors and clinical decision support; and UK guidance on evidence expectations.	These sources answer different questions about law, device classification, usability, and evidence. Their categories are not interchangeable. They cannot alone determine the hospital's review process or establish that the service, human oversight, and operational safeguards are ready locally.	Records legal applicability, device status, and evidence requirements separately from Tier 0 and the consequence tier. The hospital-specific appendix and dossier show how those determinations affect the assurance plan, workflow tests, responsible reviewers, and conditions for proceeding (**Appendices C, D, J, and L**).
Nong et al. [94]; Duke ABCDS [15]; Wisconsin CAIPA [95]; Stanford FURM [96]; Canadian action-case study [97]; and Whittaker et al. [88]	Operational precedents for registration, staged review, multidisciplinary assessment, local evaluation, monitoring, and retirement; variation in organizational design and participation.	These accounts show governance in particular settings. They do not determine how another hospital should assign authority, involve affected people, staff the work, or obtain shared support across its AI uses. Each hospital still needs a local process that identifies who is accountable for each decision and where it is recorded.	Provides a common dossier for intake, staged decisions, named signatories, approval conditions, monitoring, change, incidents, and retirement. Minimum, comprehensive, and shared-network options help hospitals assign these functions and arrange participation according to local authority and capacity (**Appendices A, B, I, and L**).
Wang et al. [5]; HAIRA [7]; PPTO [98]; and Canadian PPTO action-case study [97]	Maps existing governance approaches and gaps; describes readiness, resource constraints, and capabilities across people, process, technology, and operations.	These works describe governance components and organizational capabilities, including lifecycle practices that depend on readiness. Hospitals must still decide what their staff, infrastructure, funding, and shared support can safely sustain, and how that limits authorization of each proposed use.	Uses capability assessment to choose minimum, comprehensive, or shared-network arrangements and name owners for required functions. Before use or expansion, records support capacity, funding, and configurations needing oversight. Unresolved resource gaps can then become approval conditions or reasons to defer (**Appendices B, C, E, I, and L**).

Prior work	What it contributes	Questions left for local implementation	What this framework adds
National Academies [1]; Friedman et al. [2]; and HALO [100]	Learning from care and linking the AI lifecycle to continuing assurance and learning-health-system cycles.	Learning-health-system foundations and HALO explain learning cycles and continuing assurance. Institutions still need to decide who acts on findings, when renewed review is needed, what evidence to keep, and how to follow corrective actions through to completion.	Records monitoring findings, incidents, reassessments, and retirement lessons alongside decisions and revised safeguards. Corrective actions have owners, deadlines, and effectiveness checks; the dossier retains supporting evidence for review and restart decisions (**Appendices E, H, I, and L**).
DECIDE-AI [34], TRIPOD+AI [48], STARD-AI [109], and PROBAST+AI [110]	Reporting guidance for different evaluation designs and an appraisal tool addressing study quality, bias, and applicability.	These tools support transparent study reporting and appraisal of methods. Hospitals still need to judge whether the evidence applies to the intended population and actual workflow, whether the service is operationally ready, and what risk remains after safeguards.	Uses the relevant tool to guide evidence requests and the local-assurance report. Records applicability, uncertainty, and unresolved gaps for reviewers, then documents the decision and any additional evidence or conditions needed before proceeding (**Appendices D, I, and L**).
Matos et al. [89]	Documents fragmented fairness measures and gaps in clinical validation, uncertainty analysis, intersectional assessment, and real-world evaluation.	The review identifies weaknesses in fairness evaluation. Hospitals still need to decide what uncertainty or inequity means for the proposed use: which safeguards subgroup findings require and whether the service may proceed.	Requires subgroup assessment suited to the use, documentation of missing evidence and uncertainty, and a recorded response to inequity findings. The dossier identifies responsible reviewers, safeguards, and follow-up review. Local reviewers must justify their choice of measures and acceptable trade-offs (**Appendices D, F, I, and L**).

Local application. **Table J1** explains the sources used to design the framework. Hospitals must still determine which legal, regulatory, professional, and standards requirements apply and what evidence is needed for local operation. Tier 0 and the consequence tiers guide hospital governance decisions. They do not determine device authorization, EU legal classification, NICE categorization, or certification; each requires its own assessment.

The Joint Commission's Responsible Use of AI in Healthcare certification [104] assesses U.S. organizational practice rather than certifying individual products.

Source scope. The ISO descriptions above are based on public abstracts. Checking conformity with individual clauses requires access to the authorized standards. CHAI and AIHGle provide voluntary guidance, and Canada's strategy provides policy context. DCB0160 is mandatory within its NHS England scope and informative elsewhere. The 11 September 2026 closing date for consultation on DCB0129 and DCB0160 does not itself establish a revised standard [105]. An Ambient Voice Technology Supplier Registry listing relies on supplier self-certification and preliminary NHS England checks; it does not constitute NHS endorsement.

Using the evidence. Local use of a source requires confirmation of its current version; AI RMF 1.0 remained under revision at the 1 October 2026 source check. The NIST Playbook suggests actions, while AI 800-4 describes limitations in monitoring methods and practice. Reporting and appraisal tools help reviewers examine evidence, but cannot establish clinical utility or readiness for local use. Published implementation approaches need to be adapted to the hospital's resources and workflow. Fairness measures must also fit the intended use; they are not a universal test of acceptability.

J3 Implementation and empirical findings

Implementation studies show how governance works within staffing, budget, and workflow constraints. Nong and colleagues [94] interviewed 17 people across 13 U.S. academic medical centres and identified well-defined, emerging, and interpersonal approaches. These categories describe organizational differences, not a maturity scale. In their 2026 scoping review, Wang, Freeman, and Magrabi [5] found that 10 of 77 frameworks combined principles with an assessment method, lifecycle coverage, and oversight. Their search ended in March 2025. This finding supports examining how those elements work together. The framework and companion record combine them in a process for reviewing and overseeing each hospital AI use.

Readiness and implementation studies informed the minimum, comprehensive, and shared-network options. HAIRA [7] proposes five readiness levels across seven domains, drawing on a systematic review of healthcare AI governance frameworks; gaps in required capabilities limit progression. PPTO [98] draws on literature, experience, and key-informant interviews. Its Canadian case [97] describes implementation in one large health system. Whittaker and colleagues [88] describe governance in Aotearoa New Zealand with explicit Māori and consumer participation. Hospitals should use these readiness criteria, implementation experiences, and approaches to participation to assess their own capabilities and decide what needs to be adapted locally.

An AI output affects care through the people who interpret and act on it, which makes the workflow part of the governance assessment. Sujan and colleagues' human-factors work [81] addresses situation awareness, workload, automation bias, teamwork, relationships, training, and skill retention. Wekenborg and colleagues' toolkit [111] helps evaluate reliance and use in real healthcare settings. Drawing on this work, the framework asks hospitals to test representative tasks and handoffs under realistic time pressure, including disagreement, fallback, and emerging workarounds. The companion record captures the test conditions, findings, responsible owners, and resulting decisions.

In a 2025 simulation of AI-drafted patient-portal responses, clinicians left some errors in AI-generated drafts uncorrected and sometimes submitted drafts unchanged. [112] The study shows why hospitals must test human review in the workflow, without establishing that every drafting system is unsafe. Budzyń and colleagues' multicentre observational study, interpreted with its published correction, found an association between exposure to computer-aided colonoscopy and lower unaided adenoma detection. Its design cannot establish causation. [66][67] These studies support prospective observation of review quality, reliance, and unaided capability. Hospitals should not assume that oversight will work, or that de-skilling will necessarily occur.

Patient experience depends on the use and setting. Lawrence and colleagues' quality-improvement study of ambient documentation [83] suggests that hospitals should examine notification and opt-out preferences locally. Adoption rates alone do not show whether patients are comfortable with the service or have a good experience. The operational owner should assess how the service changes the encounter, then choose a communication approach and feasible alternative that reflect applicable law, community expectations, and significant effects on patients. Current evidence does not support one notification model for every setting.

CHARTwatch combined risk identification with communication and a clinical response pathway, so its evaluation covered a service that included data processing, alert delivery, staffing, bedside assessment, and escalation capacity. [49] Hospitals evaluating a comparable service must test that complete pathway alongside model performance. The published general-medicine study found an association with lower non-palliative in-hospital death during the intervention period, but confounding and sample-size limitations prevent attribution of that outcome to the model alone. [49]

Obermeyer and colleagues [84] showed that accurate cost prediction could still produce serious inequity when cost was used as a proxy for need. A 2026 scoping review of fairness metrics [89] also found limitations in clinical validation, uncertainty quantification, intersectional assessment, and real-world applicability. Equity review should therefore examine how the output changes decisions, who receives access or attention, and who bears the consequences of error or exclusion. Reviewers should choose measures that reveal these effects and explain their limitations. No single fairness result should determine acceptability.

Peer-reviewed studies now offer evidence on ambient documentation beyond hospital announcements. Prospective and multi-system quality-improvement reports by Ma and colleagues [113] and Olson and colleagues [114] describe modest or self-reported improvements for some clinicians, with marked variation. A rapid review [115] found only six eligible real-world studies through 2024. Evidence on safety, patient outcomes, and cost remains limited. DECIDE-AI [34] can improve reporting of early live clinical evaluation, but does not establish that approval criteria have been met. Reviewers must distinguish evidence about the underlying model from evidence about the commercial product and the service implemented locally.

Recent randomized trials show why evaluation needs to follow the complete care pathway. In LungIMPACT, AI worklist prioritization did not significantly shorten time to computed tomography (CT) or lung cancer diagnosis, with AI available in both study groups. [116] A pragmatic trial in 16 Kenyan primary care facilities found no significant reduction in its primary outcome, treatment failure within 14 days, with large language model-assisted care; both the setting and intervention differ from inpatient hospital use. [117] In a randomized outpatient trial of two ambient scribes, documentation-time effects differed between products and clinicians reported inaccuracies. [118] Hospitals should measure the intended benefit together with errors, workload, and equity for the actual product and workflow. These findings do not establish that all applications are ineffective, and a nonsignificant result does not prove equivalence.

J4 Evidence for safe and accountable use

The hospital's cybersecurity program includes responsibilities for protection, detection, response, and recovery. AI review belongs within this program, and staff should assess the proposed architecture against those responsibilities. NIST's Cybersecurity Framework 2.0 [60] provides the broader structure. The Generative AI Profile [12] addresses third-party services, connectors, content provenance, and changing model behaviour, while the adversarial machine-learning taxonomy [61] distinguishes evasion, poisoning, privacy attacks, and misuse. Cybersecurity staff can use the taxonomy to choose plausible scenarios and test the proposed controls.

WHO/Europe's maturity guidance [62] connects cybersecurity with privacy, service availability, and patient safety. AI security therefore belongs within the hospital's wider service responsibilities. Together, these sources support finding unregistered AI, reviewing system design and data flows, controlling access, and keeping logs

that allow problems to be investigated. The internal technical owner and cybersecurity reviewer should work with safety and operational teams to test plausible misuse and failure, establish the supplier's incident responsibilities, and rehearse downtime procedures. The service must also meet applicable legal and sector requirements.

Data-quality review must also establish the authority to obtain and use the data. WHO/Europe's health-data brief [58] connects quality, interoperability, sourcing, and representativeness with privacy, equity, and human rights. Canadian privacy regulators' generative-AI principles [119] call for a documented legal basis and meaningful consent when consent provides that basis. OECD guidance [59] supplies broader policy context.

Indigenous data authority needs to be interpreted for the relevant jurisdiction and community. The CARE Principles [86] address governance internationally, whereas OCAP® [87] is a First Nations framework in Canada. The relevant Nation or community determines what applies. The UK Comment on secondary data access and public trust [90] addresses its own setting and establishes no international AI or consent rule. Collection, linkage, disclosure, retention, transfer, and reuse must follow applicable law and community authority. An AI Steering Committee decision cannot provide missing legal or community authority.

Reviewers must interpret performance evidence for the intended population, purpose, and workflow, even when a study is well reported. TRIPOD+AI [48] and DECIDE-AI [34] improve reporting, and PROBAST+AI [110] helps assess prediction-model quality, risk of bias, and applicability. These tools do not demonstrate clinical utility or grant local approval. NICE [10] relates evidence expectations to function and consequence, while FDA clinical-decision-support guidance [108] addresses narrower U.S. regulatory questions that need a separate local determination.

Evaluation compares current practice with the proposed human–AI workflow and, where feasible, examines the model's contribution separately. Testing includes difficult cases and cases that differ from the data or conditions on which the model was developed. It examines uncertainty, effects on subgroups, downstream actions, and whether uncertain cases can be safely deferred for further assessment. Acceptance criteria are set before evaluation begins so they reflect the needs of the use rather than the observed performance.

AI may shift work between people and change how they learn, sometimes increasing burden or reducing autonomy despite an intended efficiency gain. It may also affect surveillance, evaluation, scheduling, hiring, and other employment decisions. OECD work on the health workforce [68], AI in the workplace [69], and algorithmic management [70] discusses skills, worker data, work intensity, accountability, health, and consultation. Operational and human-resources owners must assess the actual service with affected workers and their representatives.

Keren and colleagues [64] discuss de-skilling, mis-skilling, and never-skilling, while Ke and colleagues [65] note the absence of direct evidence for never-skilling in medical trainees and propose a staged approach to protecting competence. Hospitals should involve workers and learners in identifying essential unaided skills and assess those skills where it is safe to do so. Administrative and human-resources uses must receive a consequence tier based on the decisions they affect; a back-office purchase can still have serious consequences.

When AI changes an encounter, hospitals need to consider how patients are informed and involved, including any requirements for notification or consent. CHAI's Governance Playbooks [82] address participation and continuing feedback. WHO guidance on large multimodal models [11] supports understandable information and data protection. NHS England's ambient-scribing guidance [51] covers local evaluation, output review, and routes for questions or concerns.

We propose public transparency for registered uses and encounter-specific notice when AI materially changes an interaction, with a practical human alternative where feasible or required. Explicit consent is required whenever law, ethics, the data use, or the intervention calls for it. An acknowledgement alone does not establish

informed and voluntary consent. Joint Commission RUAIH [104] certifies hospital practice voluntarily in the United States; it neither validates products nor sets a universal consent rule.

T-CAIREM offers resources that hospitals can adapt for role-based literacy. The AI Learning Hub [71], Start Here page [72], and Living Glossary [73] provide an accessible introduction and shared terminology. For structured teaching, educators can use the International AI in Medicine Education Working Group materials [74] and curriculum framework [75] to set objectives for their audience. Broader education pages and repositories [120, 121, 122, 123, 124] provide further examples. Educators should check current versions and terms governing reuse before incorporating the materials into a local program.

T-CAIREM's Speaker Series [125] and the TAHSN AI in Healthcare Community of Practice [126] provide examples of peer learning that other hospitals can draw on when planning their own activities. The archived patient-partner workshop page [127] describes a participatory approach that can inform planning, although it is not a ready-made course. Educators must adapt cases to the approved service, assess competence for consequential tasks, and revise materials when changes in the service alter what users need to know or do.

T-CAIREM's resident curriculum combines clinical cases, short self-directed learning, facilitated discussion, and practice in communicating about AI with patients. [76] This approach can inform undergraduate education, residency, and continuing professional learning, with the adaptations for different roles described in **Section 11**. The international working group's published framework provides a complementary educational foundation. [77] Hospitals can use these resources to develop teaching, then assess whether participants can perform the tasks and decisions required in their own roles.

J5 Core requirements and local authority

These core requirements apply to the records and decisions for every AI use. Legal and professional authority remains with the relevant jurisdiction and hospital functions; the hospital-specific appendix identifies the applicable obligations and authorized decision-makers. The institutional AI registry tracks each defined use from proposal or evaluation through production, pause, and retirement. Its dossier names the clinical/end-user, operational, and internal technical owners, the executive sponsor where required, and the delegated decision-maker. Tier 0 records whether the next activity is rejected, held, or permitted to proceed. Once there is enough information to assess the proposal, it also receives a Tier 1–4 consequence classification. Permission to evaluate or assignment of a tier does not authorize live use.

Before production use, the project team must establish the need, assess realistic alternatives, and document expected value, principal hazards, data flows, and technical dependencies. Reviewers use this information to assess the complete service and identify the specialist reviews it needs. Deeper scrutiny and independent review are required where higher consequences or greater uncertainty warrant them. Evaluation must test human oversight in the intended workflow and examine effects on patients, staff, equity, and accessibility. It must also examine supplier duties and the hospital's capacity to monitor, recover, and sustain the service. Patients, communities, practising clinicians, and frontline staff contribute evidence about these effects. Indigenous authorities participate when their rights or authority are engaged. This participation informs the decision; designated authorities and specialist reviewers retain their responsibilities.

Governance approval defines the permitted scope, baseline, and conditions. Production confirmation then verifies that the actual release meets those conditions before live use. During operation, the common fields in **Appendix I** must record what is monitored, who is responsible, and what response each finding requires. Monitoring covers clinical and human effects, technical and data performance, and service operation; any area judged inapplicable requires a recorded rationale. Approved thresholds must trigger the specified response. When a pre-approved stop threshold is reached, the affected function must be suspended and the alternative workflow activated. The same dossier must be updated for incidents, use outside the approved scope, and

material changes. Restart, expansion, and material changes require the applicable scope decision and production confirmation. The dossier should make the reasons for continued use clear without forcing reviewers to reconstruct decisions from scattered files.

Hospitals should evaluate their governance program as well as individual AI services. Evaluation assesses whether the inventory and monitoring cover all relevant uses, approval conditions are completed, incidents are addressed, and retirement is verified. It examines actual benefit, the experiences of affected people, supplier performance, and audit findings. Smaller hospitals may combine compatible roles or obtain shared expertise, provided each function has a named owner and conflicts are managed. The hospital-specific appendix may change role titles, committee placement, routing of reviews and decisions, meeting arrangements, or shared-service delivery. It cannot remove a core function or leave its decisions without an accountable owner.

These core functions must be preserved when a hospital adapts the model. Each use needs registry visibility, named internal owners, a Tier 0 outcome and, when it can be assessed, a consequence tier. Decisions must include the applicable specialist and independent findings. Production use needs governance approval followed by separate production confirmation for the approved baseline and scope. Monitoring and proportionate evidence review must lead to action and fallback when required; controlled change and incident response must support authorized restart and verified retirement. Procurement safeguards, role-based literacy, and auditable records also remain necessary. A hospital may deliver a control differently only if it provides equivalent protection and records the reason, risk implications, approving authority, duration, and review trigger. Adaptations and exceptions cannot waive a core control or binding obligation.

J6 The hospital-specific appendix

The hospital-specific appendix explains how this framework operates within the hospital: who makes each decision, which authority they act under, and what records they must keep. It has a version number and change history and identifies the applicable legal, regulatory, professional, and institutional obligations it must address. **Table J2** identifies the questions the appendix must answer. Responsible reviewers must determine which laws, standards, and other requirements apply to the hospital's role and the particular AI use. The framework itself is not a source of local law.

***Table J2**. Questions for the hospital-specific appendix.*

Question for every hospital	What the hospital-specific appendix must establish
Which legal and regulatory regimes apply to the specific AI use?	The applicable jurisdiction and regulator or competent authority; any medical-device or other classification; and the hospital's role as provider, developer, deployer, or manufacturer. The appendix records required authorization, duties after a device enters the market, reporting routes and deadlines, and the consequences of local modification or use outside the device's authorized intended purpose.
What data may be collected, used, disclosed, linked, retained, transferred, or used for another purpose?	Applicable privacy and health-information law and the legal basis, including consent where relevant; data-residency and cross-border-transfer requirements; retention rules; access and correction rights; permitted secondary uses; and Indigenous or community data-governance obligations.
How must patients, users, and affected communities be informed or involved?	Requirements for notice, consent, a reasonable alternative, and public-sector transparency; accessibility and language needs; complaint routes; meaningful human review; the ability to challenge a decision; and communication after harm.

Question for every hospital	What the hospital-specific appendix must establish
Which safety, quality, and reporting processes are mandatory?	Applicable accreditation and health-service standards, clinical-safety responsibilities, and quality-assurance requirements. The appendix establishes how to report patient-safety and medical-device events, notify professional bodies or insurers, preserve evidence, and correct records.
Is the activity research, quality improvement, evaluation, or operations?	Research-ethics requirements and any conditions for innovation programs or supervised testing environments, sometimes called sandboxes. The appendix addresses publication, secondary use, participant expectations, and the decision required to move from research into operational use.
Which workforce, procurement, and institutional rules also apply?	Professional scopes and accountability, employment and workplace-surveillance law, and collective agreements; applicable accessibility, cybersecurity, public-procurement, records, and contracting rules; and requirements for supplier evidence and review, audit, and delegated authority.

A named legal or regulatory owner should maintain the appendix, recording its scope, effective date, responsible reviewers, and primary official sources for each jurisdiction covered. Version and transition details are retained so that later reviewers can establish which requirements applied to a decision. Mandatory records and reporting deadlines must be easy to find, especially when several authorities impose different duties on the same service.

The hospital's legal role requires separate assessment for each use. Terms such as manufacturer, provider, deployer, operator, health-information custodian, data controller, and research sponsor have jurisdiction-specific meanings. Buying a product, changing its intended use, or developing a system may change the hospital's role. Several regulatory regimes may apply at once. The appendix should identify the resulting obligations and state which cannot be waived through the hospital's usual process for accepting risk.

The teams whose work is affected by the appendix should review it, with patient and community input where relevant to the use and authority involved. Privacy and cybersecurity reviewers retain separate mandates, while procurement, research, legal, regulatory, labour, and audit functions contribute when their responsibilities apply. Requesters need plain-language explanations of the requirements. The appendix's owner, effective date, changes, and sources remain in a controlled record. Earlier review is required if authority changes, the service is used in a new jurisdiction or for a new purpose, a material product change affects applicability, or an incident reveals a gap.

A Canadian appendix may need to address Health Canada's April 2026 machine-learning-enabled medical-device guidance [103] alongside privacy and health-information law, human-rights and research duties, professional standards, and incident reporting.

In the European Union or European Economic Area (EEA), reviewers should check the EU AI Act [43], including current amendments and corrigenda, the current Commission timetable [106], and applicable medical-device, data-protection, national health, and public-sector law. The framework's tiers are not EU legal classifications. A U.S. appendix may require FDA device or clinical-decision-support determinations and applicable federal and state duties.

In England, DCB0160 [99] and national clinical-safety obligations may apply. NHS England set 11 September 2026 as the consultation closing date for its standards review, so reviewers must check whether revised requirements have since been issued. [105] Singapore reviewers may consider AIHGle 2.0 [8] together with

Health Sciences Authority requirements and local law. These examples identify matters to investigate in each jurisdiction; they do not mean that every instrument applies to every use.

For each material use, the dossier should explain which jurisdictional requirements apply to the exact product and purpose. It records the populations considered, the hospital's role, required authorizations, and reporting routes. It must name the responsible reviewers and identify unresolved questions, interim limits, and the event that will trigger another review. If applicability or legal authority remains unresolved, the responsible legal or regulatory owner must determine which activities may lawfully proceed and which must wait. The delegated governance authority may authorize only those permissible activities within the recorded limits. Accepting risk cannot supply missing legal authorization.

J7 Maintaining the source register

The framework owner should maintain a dated source register. For each source, the register records its authority, version, publication date, jurisdiction, and the claims and sections it supports. Each source has an assigned reviewer and a next-check date. Superseded material should remain available to explain earlier decisions, clearly identified as outdated. Review should be brought forward when regulation changes, a serious incident occurs, a new capability or setting is proposed, a link fails, or evidence suggests that a control is no longer providing its intended protection.

A proposed revision should explain why it is needed and which decisions and records it affects. Before release through the hospital's policy and document-control process, the revision is tested against representative cases. These include lower-risk and high-consequence uses, a supplier-controlled service, and a service following a serious incident or major update. The dossier should allow a reviewer to reconstruct who had authority, which evidence and affected-person perspectives informed the decision, what uncertainty remained, and which conditions were imposed. Monitoring records should also show how owners responded when changes in performance, context, law, or acceptability altered the reasons for continued use.

J8 Public investment in AI

Public investment supports AI research and computing infrastructure, alongside access to data and workforce education. The UK is funding AI-enabled scientific discovery, including drug discovery, while European Union funding supports large-scale clinical imaging pilots. [27][28] In the United States, ARPA-H's ADVOCATE program funds development and evaluation of patient-facing cardiovascular AI agents. [29] Singapore combines broad AI research investment with a separate medical foundation model initiative designed for its local populations and care pathways. [30][31] These initiatives differ in purpose and stage of development; funding commitments do not establish clinical benefit.

In Canada, AI for All includes a health-focused mission and investment in health-data infrastructure. Other announcements support the Vector Institute, AI work placements through Mitacs, and national AI literacy through Amii. [23][24, 25, 26] These programs may help hospitals build the resources and skills needed for adoption, but they do not replace assessment of an individual use. **Table J3** lists selected programs and their announced funding, distinguishing funding commitments from clinical results.

Table J3 describes selected public investments checked on 1 October 2026 and the planning questions they raise for hospitals. Currency abbreviations are CAD for Canadian dollars, GBP for pounds sterling, EUR for euros, USD for U.S. dollars, and SGD for Singapore dollars. Amounts are in their original currencies and refer to announced commitments, program funding, or calls, not audited expenditure. Some allocations form part of broader strategies, so the figures must not be added into a cross-country total. Hospitals should check current

eligibility and deadlines before pursuing an opportunity. The planning suggestions in the final column are recommendations of this framework.

***Table J3**. Selected public investments in AI research and healthcare.*

Jurisdiction and initiative	Funding status and scope	Considerations for hospital planning
Canada: AI for All, health mission, and data infrastructure (4 June 2026)	The federal strategy announced three commitments: CAD 200 million for the first AI Missions Program mission, focused on health outcomes; CAD 100 million for a Health Sector Data Space with the Canadian Institute for Health Information; and CAD 100 million to expand VITAL into five additional provinces. These commitments do not establish that all funds have been spent or that benefits have been achieved. [23]	The clinical problem, data-sharing arrangements, and prospective outcome measures must be defined before support is sought. Data access still requires lawful authority and local governance.
Ontario: Vector Institute (25 September 2026)	Ontario announced CAD 30 million for the Vector Institute to support responsible AI adoption in leading industries and develop the province's AI workforce. The funding covers research, adoption, and talent across sectors; it is not dedicated to hospital implementation. [24]	Research and technical partnerships may be considered while the hospital retains accountability for evaluation, procurement, and clinical decisions.
Canada: Mitacs AI Advantage (25 September 2026)	The federal government announced up to CAD 162 million over five years, beginning in 2026–27, for 10,000 co-funded work placements. ADOPT supports AI adoption, and AI+X supports research and innovation across disciplines, including healthcare. The allocation is not exclusively for health. [25]	Academic and industry partnerships may support implementation research and evaluation, subject to partner eligibility and co-funding requirements. Trainees should support the responsible owners without taking over their continuing operational accountability.
Canada: National AI Literacy Initiative with Amii (9 September 2026)	A CAD 13 million national partnership supports free learning streams for students, educators, and the public. Student and educator rollout had begun at the review date; the first complete public course was planned for later in 2026 through community partners. [26]	General literacy materials may be adapted where useful and supplemented with training on approved hospital tools, confidentiality, clinical limitations, verification, and escalation. General literacy alone does not demonstrate competence to use a clinical system.
Canada: AI Sovereign Compute Infrastructure Program (2026 call)	The program page specifies approximately CAD 890 million for the infrastructure build layer over seven fiscal years from 2026–27. Applications closed on 1 June 2026. Funding is expressly described as notional and remains subject to negotiation and required spending authorities. It supports general AI infrastructure rather than a health-specific allocation. [128]	Domestic computing infrastructure may support research and data stewardship. Its location alone does not resolve access control, lawful disclosure, cybersecurity, or service continuity.

Jurisdiction and initiative	Funding status and scope	Considerations for hospital planning
United Kingdom: UK Research and Innovation AI strategy (19 February 2026)	UK Research and Innovation committed GBP 1.6 billion directly to the AI sector for 2026–2030, covering research, skills, responsible AI, and infrastructure. This includes delivery of previously announced support of up to GBP 137 million for AI-enabled science, beginning with drug discovery and new treatments. The GBP 137 million is part of the headline allocation, not additional funding. [27]	Hospitals may develop multidisciplinary research and evaluation expertise through these programs, while assessing what further evidence is needed to translate early discoveries into clinical services.
European Union: AI and GenAI medical imaging pilots (2026 call)	The Digital Europe Programme call offers EUR 9 million in total for two large-scale pilots of cloud-based AI in clinical imaging. It opened on 21 April and was scheduled to close on 1 October 2026. Planned uses include flagging findings for qualified professional review and prioritizing critical cases. [28]	Hospital budgets need to cover integration, qualified review, equitable access, and evaluation of outcomes across the full imaging pathway. Receiving pilot funding does not establish diagnostic or service benefit.
United States: Advanced Research Projects Agency for Health (ARPA-H) ADVOCATE program (9 September 2026 awards)	ARPA-H announced contract awards within a four-year, USD 62.7 million program, with up to USD 33.7 million committed in year one. The work covers patient-facing cardiovascular AI agents, supervisory AI, and implementation, with an independent evaluation partner. FDA authorization is a program goal; the announcement establishes neither authorization nor effectiveness. [29]	Agents require defined action permissions, human escalation, independent evaluation, and continuing oversight. Supervisory AI also needs testing as a safety component; adding a second model does not establish that the first is safe.
Singapore: national AI research and medical foundation models (2026)	On 24 January, Singapore announced over SGD 1 billion for public research and talent under its National AI Research and Development Plan for 2025–2030. Separately, the SIMFONI initiative, supported by the Ministry of Health and launched on 9 July, develops medical foundation models for local populations, guidelines, and care pathways. The SIMFONI announcement gives no funding total and describes validation and piloting before clinical deployment. [30, 31]	Suitability for local populations, languages, standards of care, and workflows requires assessment before adoption. The national research budget must not be attributed to SIMFONI or to healthcare alone.

Appendix K Glossary

The definitions in **Table K1** provide consistent meanings for terms used in the registry, dossier, and companion record. They do not replace legal or regulatory determinations in the relevant jurisdiction. Hospitals may use local role titles as long as they preserve the responsibilities and distinctions defined here.

***Table K1**. Terms used in the governance framework and record.*

Term	Meaning in this framework
Acceptance criteria	The results a test or evaluation must achieve to pass. Before testing, the criteria specify what will be measured, the population or workload covered, the required threshold, how uncertainty will be handled, and who may accept the result. Completing a test does not mean it passed.
Action-taking or agentic system	An AI-enabled system that can use software tools or take actions beyond producing an answer. It may write to a record, send a message, place an order, schedule work, or carry out a transaction. How independently it can act varies between systems.
Adherence to approved use	Following the hospital's approved rules for the AI use in a way that can be checked. These rules cover the permitted scope, required checks and formal confirmation of outputs, what patients are told and the choices offered to them, limits on data and actions, alternative workflows, escalation, and incident reporting. The interface, permissions, audit records, training, and supervision must make the rules workable in practice.
AI dossier	The maintained record for one specific AI use. It links the evidence, reviews, conditions, decisions, changes, incidents, monitoring, and retirement record. Previous versions are retained and can be identified.
AI dossier coordinator	The person who maintains the dossier, links evidence, records decisions and conditions, and tracks review dates. The project lead or one of the three accountable owners may take this administrative role. It carries no approval authority and does not replace the governance lead's responsibility for assigning the required reviews and maintaining the recorded status of the use.
AI registry	The hospital's index of specific AI uses, with links to their dossiers. Each entry identifies the responsible owners, versions, risk tier, current status, governance and production confirmation decisions, monitoring, and changes. Statuses include Research only, Approved but not live, Routine operation, Paused, and Retired.
AI Steering Committee (or equivalent)	The multidisciplinary institutional committee that brings together the required reviews and makes or recommends specified decisions about AI use. A written mandate states which decisions it may make. A hospital may assign this work to an existing committee through formal designation.
AI system	The complete arrangement that delivers an AI-supported service. It includes the model or service and its data, settings, prompts, information sources, user interface, and connections to other systems. It also includes the users and workflow, actions based on outputs, vendor services, monitoring, and alternative arrangements when the AI cannot be used.

Term	Meaning in this framework
Approved baseline	The documented setup covered by governance approval and checked at production confirmation. It identifies the exact software or model release, settings, prompts, data connections, interfaces, workflow, safeguards, and other elements that could change behaviour or risk.
Approved limit on configuration groups	The maximum number of configuration groups the hospital agrees to support at once for one service. Approval must be justified by funded capacity to evaluate each group, educate staff for their roles, monitor and support it, respond to incidents, restore a previous approved setup when needed, and retire it safely. The proposed default of one group is a planning safeguard, not a scientifically established safety limit.
Approved operating limits	The conditions within which the AI service is approved to operate and its evidence and safeguards remain applicable. These cover purpose, population, setting, users, data, interfaces, permissions, versions, workload, and actions based on outputs. The dossier states which variations are allowed and which changes require another review.
Audit trail	Records showing what happened, when it happened, who or which system acted, the versions and permissions involved, and the result. They must support investigation of events and reconstruction of decisions or actions with significant consequences.
Classification and other predictive systems	Systems that turn inputs into a defined score, probability, category, ranking, or label. Classification assigns a category, such as whether a patient meets the criteria for a deterioration alert. Other prediction tasks may estimate a score or probability. Their outputs have a defined form, unlike freely generated content. Governance depends on the decision or action the output influences.
Clinical/end-user owner	The person accountable for whether the AI use meets clinical, professional, or end-user needs. This includes the intended purpose, human oversight, fit with the workflow, training, effects on users, and safe practice.
Comprehensive governance model	A governance arrangement with dedicated staff, continuing access to reviewers, shared infrastructure, reports covering all hospital AI uses, and more extensive evaluation than the minimum model. Both models retain every core control, including the knowledge and skills required for each role and independent review when risk or uncertainty requires it.
Configuration	The particular setup of an AI service. It includes the applicable model and software versions, prompts or templates, thresholds, data connections and information sources, interfaces, permissions, workflow, and safeguards. Recording these elements lets reviewers identify exactly what was tested and authorized.

Term	Meaning in this framework
Configuration group	A documented group of setups for the same defined AI use that can be evaluated, monitored, and supported together. Evidence must show that their allowed differences stay within approved operating limits and do not undermine the shared evaluation or safeguards. The hospital must still identify and support each setup, monitor it, investigate its incidents, restore an earlier approved setup or disable it safely, and retire it. A shared product name, site, or version number does not establish equivalence. A difference outside the approved limits requires material-change review and may require a separate group.
Control or safeguard	A technical feature, working procedure, or assigned responsibility intended to prevent, detect, or reduce harm. Examples include access restrictions, required clinical review before an action, monitoring linked to a response, and a tested alternative workflow. A proposed control must be distinguished from one that has been put in place and verified.
Corrective and preventive action (CAPA)	Work to correct a detected problem, address recurrence or similar failures, and verify that the action worked.
Credible consequence	Harm that could plausibly result from failure during intended use or reasonably foreseeable misuse. The record documents how the harm could occur and identifies uncertainty that could change the tier. The higher applicable tier is assigned provisionally while that uncertainty remains.
Data egress	Data leaving the hospital's approved internal environment or being processed outside it. This includes prompts, outputs, logs, and information derived from them.
Data flow	Where information comes from, how it moves and is processed, who can access it, and where it is stored, retained, or deleted. The record covers inputs, outputs, logs, and other information produced by the service, including movement to suppliers or other organizations.
De-skilling	Loss of an established cognitive or procedural skill through reduced practice or over-reliance on AI.
Delegated approval authority	The person or committee authorized by the hospital's written delegation to make a specified, recorded governance or production confirmation decision. The authorized decision-maker depends on the decision and risk tier and may be the AI Steering Committee, a delegated officer, Executive Leadership, or the Board or equivalent body.
Early-adopter group	The first authorized users, team, or service area. They are selected because they have a demonstrated need, represent the intended workflow, can support the service, and are willing to report problems. Their use stays within an approved scope, with defined criteria for expansion.
Early-adopter implementation	The first limited introduction into live service after governance approval and production confirmation. It covers named users or a defined service area and includes support, monitoring, an alternative workflow, authority to stop use, and a first review of operation.

Term	Meaning in this framework
Executive sponsor	The senior leader accountable for the service's purpose, expected benefit, resources, and continuing institutional support. The sponsor may accept risk remaining after controls only within their delegated authority.
Expansion decision	A recorded decision to extend the approved use to another site, specialty, population, language, or user group, or to increase volume, connections, features, actions, or independent operation. It follows review of relevant operating evidence and capacity. The expanded scope also needs production confirmation before live use.
Fallback and rollback	Fallback is an alternative way to continue work when AI is unavailable or cannot be used safely. Rollback restores an earlier approved technical setup after a change. Rollback does not by itself undo clinical decisions or transactions already made; their consequences also need review and correction. If rollback is unavailable, the hospital needs an approved, tested plan to disable the service safely and continue the work.
Focused review	A review limited to a specified change, issue, or part of the evidence. It also checks whether assumptions and controls outside that focus still hold. Unchanged items need not all be reviewed again, but any required finding affected by the change must be reconsidered.
Full review	Review of all evidence that applies to the defined use and current decision, with the depth of review matched to risk. Full review does not mean every system-specific checklist item applies.
Generative AI	AI that creates content such as text, images, audio, or code in response to an input. That content must be evaluated for the intended task and workflow. Output can sound convincing and still be wrong. Permission to generate a draft does not itself authorize the system to take further action.
Governance approval	A decision by the delegated authority that a defined AI use and the risk remaining after controls are acceptable for a specified activity, under stated conditions. Before live use, production confirmation must also verify that every release condition is met for the authorized setup and scope.
Governance lead	The person who coordinates intake, assigns required reviews to the appropriate functions, involves the necessary decision-makers, and maintains the recorded lifecycle status, including research, review, live use, restriction, pause, and retirement. The lead checks whether required records are complete and which decisions or actions must occur before others can proceed. The lead may approve a decision only under an express delegation.
Intended use	A precise account of what the AI is for, who will use it, who may be affected, and where it will operate. It identifies the inputs and outputs, decisions or actions that may follow, limits on use, and operating conditions.

Term	Meaning in this framework
Internal technical owner	The person accountable for the technical setup used in operation. Their responsibilities include connections to other systems, identities and access permissions, cybersecurity, reliability, version control, technical monitoring, and response to technical problems.
Live use or production use	Use in which the system's outputs or actions can influence actual care, service delivery, records, resource allocation, or other operational decisions. A small or early-adopter rollout is still live use. This differs from a silent evaluation whose outputs cannot guide care or operational decisions or be entered into clinical records as actionable information.
Hospital-specific appendix	The hospital's supplement to the framework, with a recorded history of versions. It states the applicable legal and institutional requirements, local role titles, decision authorities, procedures, and evidence locations. It adapts the framework to the hospital while preserving its core requirements.
Local assurance	The hospital's evidence and review that the proposed service is suitable for its intended local use and risk tier. This covers the exact technical setup, data flows, workflow, people and operating capacity, safeguards, and arrangements for continuing work if AI is unavailable or unsafe.
Material change	A change that could affect the approved purpose or population, performance, safety, equity, data handling, workflow, independent action, regulatory status, dependencies, monitoring, or remaining risk. Owners judge the change against the recorded operating limits and refer uncertainty for review before release. A change in software version number alone does not settle whether the change is material.
Meaningful human oversight	Human review that can influence what happens. The reviewer needs the skills, information, time, authority, and practical alternative needed to assess an output, disagree with it, and act in time to prevent or limit harm.
Minimum governance model	The smallest hospital arrangement that still provides every core governance function. It assigns authority and ownership, keeps the required records, and provides review proportionate to risk. It also covers governance approval, production confirmation, monitoring, response to incidents and changes, and retirement.
Never-skilling	Failure to learn an essential skill for independent work because training relies mainly on AI-supported workflows.
Operational owner	The person accountable for putting the service into practice and keeping it operating. This covers the workflow, staffing, delivery of training, support, alternative workflows, incident response, and adoption by users.
Portfolio	All AI uses for which the hospital is responsible. Reviewing them together allows the hospital to compare their benefits, risks, demands on staff and services, and unresolved conditions.

Term	Meaning in this framework
Production confirmation	A recorded decision confirming that the proposed service is ready for live use under its governance approval. It verifies the exact version and setup, scope, controls, staff readiness, monitoring, support, and alternative workflow against every applicable release condition. It is required before initial release and for the affected scope of an expansion, restart, or material release change.
Recurring governance tasks	Six kinds of work repeated as needed throughout the life of an AI use: assigning accountability and risk limits; defining the use and expected value; reviewing data and sourcing; assessing readiness; reviewing operation and responding to problems; and sharing learning and managing retirement. These describe the work to be done, rather than the seven decision stages in **Table 1** or the service's recorded status.
Reuse of evidence	Accepting current evidence from an earlier review after documenting why it still supports the claim being considered. The check must establish that the configuration, use, population, workflow, data, and controls remain equivalent for that claim. Reusing evidence does not automatically authorize a new scope of use.
Reward hacking	Pursuing a specified reward or performance score by exploiting gaps in how success is measured, rather than achieving the intended outcome. A reward is the feedback used to train or assess a system. Reward hacking is distinct from an ordinary factual error or shortcut learning.
Risk appetite, risk tolerance, and residual risk	Risk appetite is the Board's or equivalent body's broad statement of which risks the hospital is willing to consider. Risk tolerance translates that statement into practical limits on decisions others may make. Residual risk is the risk that remains for a specific use after its safeguards have been assessed.
Risk tier	A Tier 1–4 classification based on the most serious credible consequence of failure in the defined use. The assessment considers limits and human decision points only when they are implemented, enforced, and verified. Proposed or untested safeguards cannot lower the tier. Tier 0 is a separate initial screen. Risk remaining after controls is assessed separately after reviewing the evidence and safeguards.
Shared/network support model	An arrangement in which hospitals share reviewers, methods, templates, infrastructure, or learning through a network or shared service. Each institution retains authority for its own governance decisions and remains accountable.
Shortcut learning	Reliance on incidental or spurious patterns that predict labels in development data but may fail in a different setting. For example, a model may use features associated with an imaging device or documentation practice as a proxy for disease. Shortcut learning does not require data leakage or reward hacking.
Silent evaluation	An approved evaluation in which the system may use live data, but its outputs are kept from influencing actual decisions or operational records.

Term	Meaning in this framework
Specialist finding	A recorded conclusion from a reviewer responsible for a field such as privacy, cybersecurity, clinical safety, law, or research ethics. It identifies the evidence reviewed, whether requirements are met, any conditions, and activities that must wait. A committee cannot override a binding requirement within that reviewer's responsibility.
Specific AI use	A defined hospital application of AI that is registered and reviewed for a decision. The definition covers purpose, population, setting, users, workflow, data, setup, actions based on outputs, and limits. Approval applies to this specified application, not to every possible use of the product.
Stop threshold	A condition approved in advance that requires the affected function to stop immediately, an alternative workflow to begin, and escalation. Restart requires evidence that the cause and consequences have been addressed, an authorized restart decision, and production confirmation.
Streamlined review	A faster review route that the hospital may authorize for eligible Tier 1–2 uses. The use must match an approved baseline, including version, setup, data route, workflow, and controls. Staff in the relevant roles must have current competence, decision records must be complete, and the review process must have been tested with satisfactory audit findings. Ownership, required specialist findings, release checks, monitoring, incident response, change controls, and retirement duties still apply.
Tier 0 screen	A screening decision about the next proposed activity. Reject means the use is prohibited or unacceptable. Hold means essential information or a prerequisite is unresolved. Proceed permits the specified review or evaluation to begin. The decision identifies required specialist work and the steps that must be completed before other work can proceed. It does not authorize live use.
Vendor POC authorization	Permission for a time-limited purchasing evaluation under a recorded plan. The plan specifies the products and setups, question to be answered, data and test environment, criteria, responsible people, costs, dates, and disposal arrangements. Outputs cannot be relied on in operation. This permission does not itself grant data access or replace required specialist clearances.
Vendor POC results review	The procurement record comparing the exact vendor product and setup with selection criteria set before testing. It records whether to stop, seek more information, change the specification, compare another option, negotiate conditions, or continue procurement. It grants neither governance approval nor production confirmation.
Vendor proof of concept (POC)	A time-limited evaluation of a vendor product or setup before committing to a commercial arrangement or choosing the configuration. A separate POC is required when existing evidence or an equivalent authorized evaluation cannot resolve important questions. If it is omitted, the procurement owner and responsible governance authority must record the evidence and reasons, subject to local policy. A POC does not replace local assurance or permit reliance on its outputs in operation.

Term	Meaning in this framework
Workflow or use drift	A change in how people use the AI compared with the approved workflow. This includes how they check outputs, rely on them, bypass expected steps, or act on results, even when the model itself has not changed.

Appendix L Hospital AI Governance Intake and Lifecycle Record

This record documents one defined AI use throughout its time in the hospital. The requester describes the proposal; designated owners and reviewers then keep the evidence, decisions, conditions, and updates current. The AI registry is the hospital's inventory of uses. The dossier is the case file for this use, containing this form and links to supporting evidence. When several records rely on the same evidence, each can link to one stored copy. The main framework explains the requirements, and **Appendix I** defines the fields shared across the records.

Hospitals may adapt labels, review assignments, the placement of roles, or evidence storage. Each local change must preserve the core requirements in **Appendix J5** and the field meanings in **Appendix I**. The change record explains the reason, how equivalent protection is maintained, who approved it, its effective date and expiry, and what would trigger review. A local change cannot waive a core safeguard or binding obligation.

This record links to the hospital policies and decisions that govern the use, including the hospital's risk appetite and delegated decision powers, the registry and approved-services list, committee decisions, and oversight records for its AI services. It also links to the relevant contracts, literacy programs, lessons from incidents, and approved hospital-specific appendix. These links help reviewers find the applicable policies, earlier decisions, and responsible decision-makers. **Table L1** maps the sections of this record to the main framework and related appendices. Within this record, section numbers refer to **Appendix L** unless otherwise stated.

***Table L1**. Framework requirements and companion record sections.*

Intake record sections	Main framework sections	Related appendix
Working record and **Sections 1–3**	**Sections 6–7**: registry visibility, owners, problem, value, and intended use	**Appendix I**: core fields
Section 4	**Section 9**: Tier 0, Tier 1–4, depth of evidence review, and routing for specialist review	**Appendix D**: evidence by tier
Sections 5, 5.2, and 5.3	**Section 8** and **Section 7.6**: procurement and moving research into operational use	**Appendix G**: procurement, data leaving approved systems, vendor proof of concept (POC)
Sections 6–7	**Sections 8–11**: privacy, cybersecurity, safety, people, evidence, and system features	**Appendix F**: system-specific checks
Section 8	**Sections 10, 11, and 12**: readiness, monitoring, change, incidents, and retirement	**Appendices E and H**
Sections 9–10 and reusable records	**Sections 7, 8, 10**, and local authority recorded in the version-controlled appendix	**Appendices H–J**

How to use this record

This form follows one proposed AI use through review and, if approved, into hospital practice. It belongs in the dossier and links to evidence held in approved systems. It describes the product and release, task, setting, affected people, and permitted actions. It identifies where a person must make a decision before anyone relies on the output. The requester completes **Sections 1–3**. The governance lead coordinates screening and reviews, records decisions, and keeps the status current. The dossier coordinator maintains the form and evidence links. A materially different task, action, patient or user population, or workflow requires an approved change or a separate record.

A configuration is the complete setup being reviewed: the product and release, settings, data routes, workflow, safeguards, and permitted uses. Once approved, this setup becomes the approved baseline against which later changes are checked. Variants of the same defined use may form a configuration group and share governance and monitoring arrangements, but only when a documented comparison shows that they are equivalent for those purposes. The group record specifies the permitted differences and approved operating limits for population, context, interfaces, data, permissions, versions, and workload. A different site or software version does not automatically require a new group; the comparison must establish whether the existing arrangements still apply.

Sections 1–3 require enough information to start review, with anything still unknown identified and a named person responsible for supplying it. Submitting the form does not establish readiness for approval. In **Section 4**, the governance lead identifies the reviews and evidence needed before the next activity. **Section 5** assigns responsibilities for building, buying, or partnering. **Section 5.2** covers a time-limited vendor proof of concept (POC) to answer a purchase or configuration question. **Section 5.3** covers a proposal to move research into operational use. **Section 7.3** records local assurance—the hospital's checks of the proposed setup, data routes, users, and work process. The checks required depend on the use and the decision requested.

Section 10 separates two decisions. Governance approval defines what is allowed and under which conditions; production confirmation checks the actual release and verifies that every condition due before live use has been met. Expansion needs approval of the increased scope followed by production confirmation. The first authorized operating scope, called early-adopter implementation, requires a setting with a demonstrated need, representative work, capable local leadership, and the ability to identify and contain problems. Early-adopter implementation is a service status, as are routine operation, restriction, and pause. It is not a review stage or governance task. A vendor POC is a procurement evaluation with stated data and access limits. It does not create an operating status or permit outputs to guide care or hospital operations.

Before a decision, the reviewer identifies the requirements that apply and when each item of evidence is due. Evidence links point to existing records where possible. Met means the reviewer has confirmed from the evidence that the requirement is satisfied for this decision; the verifier and date are recorded. An unresolved requirement remains Open, classified as Pending, Failed, Unknown, or Expired. Its record identifies which decision must wait, the action needed, its owner and deadline, how completion will be checked, and how an overdue response will be escalated. Evidence needed only for a later decision stays Open—Pending and does not delay earlier work unless that work depends on it.

N/A applies only when the reviewer accepts and records why a requirement does not apply. N/A cannot waive a mandatory safeguard or required authorization. If a row contains several requirements, every applicable part must be met before the row is marked Met. Ready means all requirements needed for the decision are Met or have an accepted N/A reason. Ready with required actions means the remaining work can safely wait until after this decision; each action needs an owner and deadline. Not ready means a requirement needed now remains unresolved.

Patient-level information, passwords, security secrets, and source code belong in the appropriate controlled systems and must not appear in this form.

This record supports the hospital's existing governance processes. **Appendix I** defines the registry and dossier fields. Evidence remains in approved systems with controlled access and identifiable versions. The linked-record table identifies the version considered in each decision. The hospital-specific appendix records decision authorities and any

adaptations to the framework. Its version number and change history must remain current and explain each adaptation without removing core safeguards or mandatory duties.

The record begins with the need and proposed use. Tier 0 then determines whether the next activity may proceed to risk tiering and evidence review. Governance approval sets the permitted use and conditions, and production confirmation is required before live use. During operation, monitoring and review findings inform decisions on changes, expansion, incident response, and restart. A changed scope or restart requires the applicable approval and production confirmation before release. Withdrawal or replacement requires completion of the retirement process.

Working record

The governance lead completes the Working record with the requester and keeps it current during review. It identifies the accountable hospital, the rules governing the review, and the people authorized to decide. The record states the exact permission requested and the reason for any fixed deadline. It separates the service's current status from the decision stage, so reviewers can distinguish what is happening now from what they are being asked to approve. **Tables L2–L5** record the hospital and review authority, current status, requested decision, and submission details.

***Table L2**. Accountable hospital and review authority.*

Hospital or health system The record identifies the accountable hospital and any participating sites.	
Jurisdiction and governing requirements The record identifies the jurisdiction and the principal law, policy, or local standard governing this review.	
AI use name and tracking references The record includes the name of the AI use, hospital tracking number, and governance case ID. These references link this form to the registry, dossier, and governance software or controlled manual log.	
Submitting team The record names the team making the request.	
Who can approve this request? The record names the people who can give governance approval and production confirmation. It states which tiers and decisions each may approve. It links to the policy or delegation that grants that authority.	
Hospital policies and decision limits The record links to the current AI policy, Tier 0 criteria, delegated decision powers, and current version of the hospital-specific appendix. It also links to the hospital's risk appetite and tolerance: the risks it is willing to accept and the limits it sets.	

Version and approval of this record The record provides the version, person responsible for maintaining this record, approval date, effective date, and next review date.	
Local changes to the framework Each change has a recorded reason, explanation of equivalent protection, named approver, and expiry. Evidence demonstrates that core safeguards and mandatory authorizations remain in place.	

***Table L3**. Current service status.*

Current service status The selection describes the service now. Details include when that status began and the current decision stage.	☐ Research only ☐ Intake ☐ Local assurance ☐ Approved but not live ☐ Early-adopter implementation ☐ Routine operation ☐ Restricted ☐ Paused ☐ Retired Details: status start date; current decision stage; decision reference.

Decision stage means the part of the review process shown in **Table 1**: 1 Intake and triage; 2 Problem and value; 3 How the service will be provided and prepared; 4 Evidence and governance decision; 5 Production confirmation before live use; 6 Continued use and changes to approval; or 7 Retirement or replacement. Lifecycle status means the service's current condition. For example, a service can remain in routine operation while a proposed expansion is reviewed at stage 6. Expansion is a decision at that stage, not a lifecycle status.

***Table L4**. Requested decision, scope, and timing.*

Decision requested The record identifies one requested decision.	☐ Vendor POC ☐ Governance approval ☐ Production confirmation ☐ Expansion ☐ Continue ☐ Change ☐ Restrict or pause ☐ Restart ☐ Retire ☐ Other For production confirmation, the record identifies whether the request covers initial release, a material release change, expansion, or restart.
Exact permission requested The record describes what would become newly permitted. It specifies the service and activities covered, users, setting, actions, dates, and limits.	
Timing and fixed deadlines The record provides the requested decision date and earliest date of operation. It explains why any deadline cannot move.	

***Table L5**. Request type and submission date.*

Request type All applicable request types and the submission date are recorded.	☐ New use ☐ Existing live use ☐ Purchase/contract ☐ Vendor POC ☐ Research transition ☐ Expansion ☐ Material change ☐ Scheduled review ☐ Incident ☐ Retirement Details:

For each item, **Table L6** records the approved file or record location, version or date, owner, and status, using the **Appendix I** fields. Each entry links to the stored evidence so the same file does not need to be copied into this record. Earlier versions must remain clearly identified when a record changes so reviewers can reconstruct the evidence available at each decision.

***Table L6**. Supporting records and evidence locations.*

Record or supporting file	Identifier / location / version or date	Owner / status
AI registry entry		
AI dossier		
Governance workflow/case record and audit trail of actions and changes		
Review-route qualification and capacity record		
Standard operating procedure (SOP) / user instructions		
Map of data flows and technical system design		
Evidence / local-assurance report		
Record of possible failures, harms, and safeguards		
Monitoring plan / dashboard		
Training and competency record		
Patient, user, or public information		
Contract / vendor / responsibility record		
Change and release log		

Record or supporting file	Identifier / location / version or date	Owner / status
Incident, restart, and retirement record		

1 Use, ownership, and status

The requesting team describes the product and functions proposed for the first operating scope, distinguishing them from features that will remain disabled or may be considered later. The team identifies who will fund and support the work, discloses any current use and whether it needs to stop or be restricted, and names the accountable hospital owners. Its account explains how affected staff and communities have contributed or, if involvement has not started, who will arrange it. **Tables L7–L12** record the proposed scope, funding, affected areas, enabled system functions, sourcing route, research status, and current use or temporary controls.

Table L7*. Proposed use, operating scope, and funding.*

Plain-language summary The record explains in one or two sentences what is proposed, who would use it, and how the current work would change.	
First proposed use The record specifies the initial sites or services, authorized roles, expected volume, and period of support. An exploratory proposal identifies what has not yet been decided.	
Why this proposed use is appropriate The record explains the demonstrated need and how the proposed work represents the intended use. Evidence demonstrates that the team can report and contain problems. It states the maximum extent of use requested now; any future expansion is described separately.	
Implementation funding The record lists the one-time work needed to introduce the service. It provides its funding source, budget owner, committed funding or capacity, and availability date.	

Ongoing operating funding The record names the funding owner, funded period, and renewal date. Funding covers licensing or hosting, monitoring, training, support, incident response, continuity, and exit.	
Funding risks The record explains any gap, dependency, expiry, or renewal risk that could affect this decision. It states what will happen if funding ends.	
Vendor or partner support The record discloses discounts, free licences or credits, grants, vendor-funded evaluation or integration, and other in-kind support. It includes the conditions and expiry. It explains how conflicts of interest will be managed.	

***Table L8**. Areas affected by the use.*

Area of use The record identifies every area affected.	☐ Care ☐ Patient service ☐ Operations ☐ Administration ☐ Workforce/human resources ☐ Research ☐ Education ☐ Public health Details:

***Table L9**. System functions, authority, and configuration.*

Model or system type The record identifies every applicable model or system type.	☐ Prediction or classification (discriminative) ☐ Content generation (generative) ☐ Task execution or coordination (agentic/workflow orchestration) ☐ More than one type ☐ Other/unknown Details identify any ambient scribe functions, such as recording an encounter or drafting documentation, and models managed by an outside provider or selected by automated routing. This allows reviewers to assign the relevant checks for each system capability. Details:
Function The record identifies every enabled or planned function.	☐ Capture/transcribe ☐ Retrieve ☐ Summarize ☐ Draft content ☐ Extract structured data ☐ Score/classify ☐ Recommend ☐ Prioritize ☐ Communicate ☐ Create/modify record ☐ Initiate/carry out action Details:

Actions the system may take The selection shows the highest level of action authority currently enabled or requested. The record identifies separately any features planned for later or currently disabled.	☐ Information only ☐ Draft requires full human review ☐ Recommendation requires a human decision ☐ Action is prepared but requires confirmation ☐ Action is carried out under direct human supervision ☐ Action within defined limits can be reversed and does not require case-by-case confirmation ☐ Time-critical or difficult-to-reverse action does not require case-by-case confirmation Details:
How the system can change The selection includes every way its behaviour can change.	☐ Fixed version ☐ Vendor-scheduled update ☐ Local retraining or fine-tuning ☐ Prompt or template change ☐ Changes to information retrieval or reference content ☐ Automated routing between models ☐ Learning during operation (online/continuous learning) ☐ Unknown Who can test, defer, reject, approve, or reverse each change?
Integration The record identifies every applicable integration.	☐ Stand-alone ☐ Embedded in the electronic health record (EHR) ☐ Connected through a software interface (API) ☐ Integrated with a device ☐ Agent connected to several systems or tools ☐ Vendor-managed service Details:
Product, configuration, and updates The record identifies the product, model, provider, and release. It lists enabled and disabled features, any other models it depends on, and any routing between models. It provides the update frequency or policy and identifies who receives change notices.	

***Table L10**. Proposed sourcing route.*

Sourcing route The selected route is the main way the service will be provided. The record identifies each component that uses a different route.	☐ Build through hospital-led development ☐ Buy/configure, including a managed service ☐ Partner/co-develop ☐ Mixed components—component map required ☐ Undecided Provider, product/model, version, hosting/model provider, and component details:

***Table L11**. Research status and operational boundary.*

Research status Where research applies, the record includes the ethics status and reference. The record explains how research outputs are prevented from guiding routine care or operations outside the authorized research protocol.	☐ Not research ☐ Research only ☐ Transition requested Details:

***Table L12**. Current use and temporary controls.*

Current-use status Is this AI already in use? Available responses: No, Yes, or Not known.	☐ No ☐ Yes ☐ Not known
Details of current use For Yes or Not known, the record describes what is known about the tool, users, data, access, start date, approval, and registration. Uncertainty and immediate risks require prompt disclosure. Stopping or restricting an unsafe use must not wait for complete information.	

Temporary controls The record states what will be restricted or suspended and which safe alternative will be used. It names the person authorizing this and the expiry date. Continuing temporarily requires separate, time-limited approval and equivalent protection. Mandatory duties still apply.	

Named people within the hospital remain accountable for this use, even when a vendor or partner performs much of the work. The clinical/end-user owner is accountable for appropriate use and outcomes. The internal technical owner is accountable for the local technical service. The operational owner is accountable for workflow, support, and sustainability. The record distinguishes the tasks others will perform from the decisions and checks each hospital owner retains. One person may hold several roles only if local policy allows it and they have the time and capability to fulfil each responsibility. **Tables L13–L15** record these owners, staff and community involvement, and any conflicts of interest and their management.

Table L13*. Named owners and supporting contacts.*

Role	Name, title, unit, contact	Responsibility accepted and date
Executive sponsor, if required		
Clinical/end-user owner		
Internal technical owner		
Operational owner		
Dossier coordinator Maintains the records and evidence links. This responsibility does not grant approval authority.		
Vendor contact accountable for technical work, if used		
Project Working Group lead and members The record identifies the governance lead separately if another person holds that role.		
Vendor-management, procurement, or contracting contact, if used		

***Table L14**. Staff and community involvement.*

Staff involvement The record names the affected roles and those who participated. It includes their recommendations and the response or next step. This includes administrative, human-resources, and back-office staff when their work is affected.	
Patient, family, community, and Indigenous involvement The record identifies the relevant people or representative bodies. It explains how their input was addressed. If involvement has not started, the record gives the plan and owner. The applicable Indigenous-led process must be followed where required.	

***Table L15**. Conflicts of interest and their management.*

Conflicts of interest A Present or Not known finding includes a description of the relationship, the decisions it could influence, and how it will be managed.	☐ None known ☐ Present ☐ Not known Details:

2 Problem and value

In **Table L16**, the requesting team and affected service leads record the problem and a dated measure of current performance. Their comparison of the proposed AI use with current practice and at least one non-AI alternative covers benefit, risk or burden, effort, and cost, including effects on people who could face extra work or reduced access. Any unanswered question that could change the choice needs a named person responsible for resolving it. The record describes the safe alternative the hospital will use if AI brings no worthwhile benefit or becomes unavailable.

***Table L16**. Problem, alternatives, and expected value.*

Problem and current performance The record explains who is affected and where the current process falls short. It includes a dated baseline showing current volume, time, errors, or another relevant measure. It provides its data source. If no baseline exists, the record names who will collect it and by when.	
Options compared The assessment compares the current process, at least one non-AI option, and the proposed AI use in terms of expected benefit, main risk or burden, effort, and cost.	
Expected benefit The record names the outcome to improve, its current value, the smallest worthwhile improvement, the assessment period, and the accountable owner.	
Added burden or disadvantage The record identifies who may face more work or cost, reduced access, or loss of a skill. It states how the burden will be measured and the highest acceptable level.	

Questions that could change the decision Each important uncertainty has a record of the information needed, who will obtain it, and the deadline.	
If the use adds no value or is unavailable The record describes the safe alternative and identifies who decides whether to stop the AI use or return to that alternative.	
Organizational objective(s) ☐ Patient outcomes ☐ Patient experience ☐ Clinician/workforce experience ☐ Safety/quality ☐ Equity/access ☐ Operational efficiency ☐ Economic efficiency/resource reallocation ☐ Research/education ☐ Public health ☐ Other	

Value decision: ☐ AI option preferred ☐ Non-AI option preferred ☐ More information needed—hold
Decision owner/date:

3 Intended use and boundaries

The requesting team, clinical/end-user owner, and operational owner record the agreed intended use and its limits in **Table L17**. This ensures that every reviewer assesses the same proposal. The account describes what information enters the system, what it produces, and what happens next. It shows where a person must review the result and whether they have the information, time, and authority to disagree. It specifies what is permitted, what is prohibited, what affected people will be told, and what choices they have. It also states when the service must stop, what will replace it, and which evidence is needed for the use and decision requested now.

***Table L17**. Intended use and operating boundaries.*

Intended use The record describes how the user performs the task, the information the system needs, the output it produces, and the decision or action that follows. It identifies the setting and people affected.	
Users, affected people, setting, and scale The record specifies sites, shifts, authorized roles, populations, exclusions, and expected volume, including any group for whom performance or access may differ.	
Inputs The record lists each input needed for the decision and where it comes from. It explains what happens if it is missing or wrong.	
Output and next action The record describes who receives the output, when it arrives, what action is permitted, and where the output is saved or sent.	

Work process and response time The record explains what starts the process, where the output appears, and who must respond. It states the expected response time, next action, escalation procedure, and what happens if no output arrives.	
AI involvement and human review The record specifies which of the framework's five AI involvement arrangements applies to each task, including tasks that remain human-only. It identifies who reviews or supervises the work and what they must check. Evidence demonstrates that they have enough source information, time, competence, and authority to disagree or intervene. It describes how people review, override, and escalate results, and how these actions are recorded.	
Permitted and prohibited uses The record describes the limits for users, population, setting, data, version or configuration, workflow, actions, and scale. The same categories specify what is allowed and what is prohibited. It explains how prohibited use is prevented.	
Information and choice The record states what people will be told, when, and by whom. It explains the consent or other basis for the activity and any option to refuse, opt out, or use a non-AI route.	

Questions, corrections, and complaints The record provides the contact and expected response time. It identifies who can review or correct the issue and what protection is available while it is being considered.	
When to stop and what to use instead The record lists the events or thresholds that require suspension. It identifies who must act, how quickly, and which safe alternative begins. Reaching a stop threshold requires that response without waiting for further permission.	
Changes to existing permission The record states what new permission is requested for data, environment, users, sites, workflow, outputs, actions, integrations, volume, or duration. It links to the Working record entry that defines the proposed use and its limits.	
Evidence needed for this decision The record links to the local-assurance plan for this version, proposed use, and limits. The plan distinguishes evidence needed now from evidence needed only for later decisions. It covers data flows, workflow, affected groups, possible failures, human checks, safe alternatives, and scale.	

Limits of the first implementation The record identifies later sites, specialties, populations, languages, users, volumes, integrations, features, or changes in autonomy being considered. It explains why evidence from the first implementation may not apply to those changes.	

I confirm that **Sections 1–3** are accurate to the best of my knowledge. I have disclosed any current unapproved or unregistered use, important uncertainty, dependency, conflict of interest, or serious concern. I understand that submitting this record does not authorize a purchase, access to live data, live testing, or implementation.

The requester signs **Table L18** to confirm the attestation above.

***Table L18**. Requester attestation sign-off.*

Role	Name	Signature or electronic confirmation	Date
Requester			

4 Screening, tiering, and review route

The governance lead and named owners first screen the requested activity. If it may proceed, they assign a provisional tier based on the consequences of failure and arrange specialist reviews. Each review has a named reviewer, a defined scope, and an identified decision that requires its conclusion. Evidence due later may remain pending; it has not yet been accepted. Work beyond planning or contract discussion requires a separate recorded authorization specifying its limits. Uncertainty about a prerequisite must be resolved before the decision that depends on it.

Tier 0 determines whether the next activity may proceed; consequence tiering follows. Reject applies to a prohibited or otherwise unacceptable use. Hold applies when essential information is missing or a barrier to the requested activity remains unresolved. Proceed applies only to the review or evaluation described in the decision; it does not permit live use. A No or Not known answer does not automatically require rejection, but the decision needs an explanation. A question that does not apply receives a check in the Yes or accepted N/A column, with the reason accepted by the reviewer recorded. Any unapproved use already under way must be stopped or restricted before review continues. **Table L19** records the answers; **Table L20** records the decision and permitted activity.

***Table L19**. Tier 0 screen for the next activity.*

Tier 0 question	Yes or accepted N/A	No or not known	Reason, affected activity, action, owner, and date
Law and policy Can the next proposed activity, including its data use, comply with law and policy? The record identifies any question requiring specialist advice.	☐	☐	
Purpose, owners, and stopping Are the proposed use, accountable internal hospital owners, person authorized to stop it, and safe alternative clear enough to assess the next activity?	☐	☐	
Privacy and cybersecurity Are the required reviews identified and scheduled before the data or system access they govern?	☐	☐	

Tier 0 question	Yes or accepted N/A	No or not known	Reason, affected activity, action, owner, and date
Required approvals Have the necessary specialist and external authorities been identified? The record identifies who will obtain each finding or approval and the decision stage by which it is needed. Approvals needed only later do not all have to be complete at intake.	☐	☐	
Unresolved requirements Have barriers to the activity requested now been resolved? Each later-stage barrier has a named owner and an identified decision that it blocks.	☐	☐	
Safety, rights, equity, and accessibility Can reviewers assess plausible serious harm, inequity, inaccessibility, and culturally unsafe effects? Can the next activity be carried out safely within its stated limits?	☐	☐	

Tier 0 question	Yes or accepted N/A	No or not known	Reason, affected activity, action, owner, and date
Indigenous governance Where Indigenous people, communities, or data may be affected, has the applicable Indigenous-led process been identified? Is it scheduled before the activity that requires its authority?	☐	☐	
Oversight and supplier cooperation Can the hospital obtain the evidence, logs, change notices, incident support, and ability to exit that it needs for oversight? Have the requirements for each later activity been identified?	☐	☐	
Current unapproved use Has any unapproved or unregistered use been stopped or restricted? The record links to any separate, time-limited permission to continue. If no such use was identified, the answer is Yes, with that finding recorded.	☐	☐	

Tier 0 question	Yes or accepted N/A	No or not known	Reason, affected activity, action, owner, and date
Benefit assessment Are the proposed benefit and decision owner identified? Have the baseline and minimum worthwhile improvement been recorded, or assigned for completion before the relevant evaluation?	☐	☐	
Criteria for accepting results Has someone been assigned to approve the safety, performance, and equity criteria before results are reviewed and before the relevant activity starts? Criteria for an activity starting now must already be approved.	☐	☐	
Human review and investigation Is there a workable plan for the required human review and for tracing what the system produced or did? The record identifies the evidence needed to verify that plan before anyone relies on the system.	☐	☐	

Tier 0 question	Yes or accepted N/A	No or not known	Reason, affected activity, action, owner, and date
Permission to access data and systems Is the access needed now authorized? The record identifies any additional data routes or secondary uses that need separate clearance before a later activity.	☐	☐	

***Table L20**. Tier 0 decision and permitted activity.*

Screening decision The recorded decision is Reject, Hold, or Proceed, with reasons and the exact activity that may or may not take place. A Hold decision identifies the evidence or change needed before the affected activity resumes. Any action to stop or restrict current use is recorded separately.	☐ Proceed to tiering and specified review ☐ Hold for information or redesign ☐ Reject Current-use containment: ☐ Not applicable ☐ Stopped ☐ Restricted under separate authority Details: allowed activity; limits; owner; date; expiry; requirements that must be met before each open condition can be closed.

4.1 Assign Tiers 1 to 4

After a Proceed decision, reviewers assign Tier 1, 2, 3, or 4 using the definitions in **Table L21** and questions in **Table L22**. Their assessment considers harm that could plausibly result from intended use or foreseeable misuse and explains the supporting evidence and uncertainty. It reflects what the system is permitted to do and which limits are enforced in practice; the product name alone cannot establish a tier. After the preliminary checks, reviewers work down the tiering questions from Tier 4 to Tier 1. The first Yes answer to a tiering question identifies the provisional tier. A limit can rule out a more serious failure only if it has been implemented, enforced, and verified. Planned or untested safeguards cannot lower the tier. Human review counts as a limit only when the reviewer has the source information, time, competence, and authority to disagree before action occurs.

***Table L21**. Consequence tiers and distinguishing features.*

Tier	When it fits	Example
Tier 1 — low	Only minor, readily corrected inconvenience is credible. The use cannot affect care, clinically significant access, patient communication, rights, employment, resource allocation, or the official clinical record.	Non-sensitive internal drafting checked before use.

Tier	When it fits	Example
Tier 2 — moderate	Failure could cause meaningful but limited disruption, delay, workload, or a correctable privacy, fairness, or service concern. Serious harm is not a credible consequence within the defined use. The problem can be corrected in time.	Scheduling or capacity support within defined limits that does not determine care or clinically significant access.
Tier 3 — high	Output can materially influence care, access, employment, rights, resource allocation, prioritization, patient communication, an official record, or another consequential action. Meaningful human review occurs before action.	Clinician-reviewed deterioration alert, signed scribe note, or human-approved coordination action.
Tier 4 — very high	The system can drive or carry out a time-critical or similarly consequential decision or action without meaningful approval by an authorized person for that case. Tier 4 also applies if serious harm can occur before meaningful review.	Autonomous time-critical triage, treatment, denial of clinically significant access, or another action capable of serious harm.

***Table L22**. Risk-tier decision checklist.*

Order	Tiering question	Result
1	Is the proposed use clear enough to identify its users, population, setting, data, output, next action, human decision authority, and scale? A No answer requires clarification before a tier is assigned.	☐ Yes ☐ No
2	Does the Tier 0 decision allow this tiering and review? Tiering and review require a Proceed decision for that activity. Reviews needed only at a later stage do not by themselves prevent tiering.	☐ Yes ☐ No/unknown
3	Can the system drive or carry out a time-critical or similarly consequential decision or action without meaningful approval by an authorized person for that case? Or can serious harm occur before meaningful review? If either answer is Yes, Tier 4.	☐ Yes → Tier 4 ☐ No
4	Can the output materially influence care, access, employment, rights, resources, prioritization, patient communication, an official record, or another consequential action? Does meaningful human review occur before action? If both answers are Yes, Tier 3.	☐ Yes → Tier 3 ☐ No

Order	Tiering question	Result
5	Could failure cause meaningful but limited disruption, delay, workload, or a correctable privacy, fairness, or service concern? Can the problem be corrected in time, with no credible serious harm within the defined use? If both answers are Yes, Tier 2.	☐ Yes → Tier 2 ☐ No
6	Are all credible consequences minor and readily corrected? Is there no material effect on care, access, rights, employment, resource allocation, patient communication, or an official clinical record? Does a person check before anyone relies on the output? If all answers are Yes, Tier 1.	☐ Yes → Tier 1 ☐ No
7	Does the proposal combine substantially different uses, or are important facts still unknown? Separate uses receive individual classifications. The higher provisional tier remains until evidence, or an effective separation of uses, rules out the more serious consequence.	☐ Done ☐ Action needed
8	Has a tested technical or workflow limit ruled out a failure with more serious consequences within the approved use? The record includes evidence that the limit is enforced and works. Training, policy, monitoring, or an expectation of careful use alone cannot lower the tier.	☐ Recorded ☐ N/A

Table L23 explains how the most serious credible harm could occur, tracing the steps from a system failure or foreseeable misuse to its effects on people. The explanation includes the supporting evidence and uncertainty; a likelihood label alone is insufficient. Where relevant, the rationale explains why the next higher tier does not apply. The record includes any disagreement and the changes that would require the tier to be reconsidered. After reviewing the safeguards and evidence, reviewers assess the risk that remains separately from the consequence tier.

Table L23*. Tier rationale, safeguards, and reassessment triggers.*

Question and information to provide	Response or evidence reference
Exact use and system authority Who receives what output, and where? What can the system or user draft, advise, prioritize, record, communicate, recommend, or carry out?	

Question and information to provide	Response or evidence reference
Most serious credible failure The record explains what would happen in the actual workflow if output were wrong, missing, late, unavailable, or misused. This includes foreseeable changes or failures that could defeat an approved limit.	
People, urgency, vulnerability, and scale Who could be affected, and how soon could harm occur? Which groups might be more vulnerable? The record estimates the largest credible number of people exposed before the problem is noticed.	
Detection and reversibility Would someone notice the problem before harm occurred? Could the action, communication, access decision, or official record be fully corrected in time?	
Human authority Who reviews or approves? What source information and time do they have, and can they disagree? How is approval recorded?	
Data, rights, and integration Does the use involve sensitive data, an external service or provider processing data, patient-facing output, or an official record? Does it affect rights, employment, or resource allocation, or involve writing back to a record or system, connected tools, or an external action?	
Enforced limits on use Which technical or workflow limit prevents the more serious consequence? Evidence demonstrates that it is implemented and enforced and explains how it was tested.	
Provisional tier and reason The record states Tier 1, 2, 3, or 4. Where relevant, the rationale explains why the next higher tier does not apply. The higher provisional tier remains while an unresolved fact could change that judgment.	
Changes that require a new tier assessment The record identifies changes that would require the tier to be reconsidered: purpose, user, population, site, model, data, threshold, prompt, template, integration, output, action, human review, autonomy, supplier, or scale.	

4.2 Assign the required reviews

For each specialist review, the governance lead records whether it is needed, who will perform it, what it covers, and which decision requires the conclusion. Not required needs an accepted reason. An Unclear finding needs a named person and deadline for resolving the question. Reuse of earlier evidence requires confirmation that the use, configuration, data routes, setting, and assumptions still match and that the finding is current. A focused review may examine one change or question while confirming that the rest of the assessment remains valid. An unresolved review requirement blocks the decision that depends on it; unrelated preparation may continue. **Tables L24–L27** record review assignments, permitted activity during review, the governance lead's sign-off, and whether independent review is needed.

A locally approved streamlined pathway is available only after Tier 0 permits review and the use is clearly Tier 1 or Tier 2. The use must match a documented approved baseline, have complete standard records, and have current evidence of competence for the people performing each role. It must have no material effect on care, clinically significant access, rights, employment, resource allocation, patient communication, an official record, or a safety-critical workflow. An unresolved specialist concern, blocking action, serious incident, or material supplier issue prevents use of this pathway.

Before the streamlined pathway opens, the AI Steering Committee must approve its eligibility rules and audit plan. The plan defines which Tier 1 and Tier 2 decisions qualify, how they will be grouped and sampled, and why the sample is sufficient. It also defines a material error, names the auditor, sets the audit frequency, and states the acceptable result. Testing against completed full-review decisions provides a baseline audit; the pathway may open only if that audit meets the criterion agreed in advance. Audits continue at the approved frequency, with results reported by tier and use group. The audit plan specifies the findings that require the named authority to suspend the pathway. Reopening requires verified completion of corrective actions, current role competency, and an audit result that meets the agreed reactivation criterion. Tier 3 and Tier 4 remain excluded. **Section 7.3** links to the central program records and current pathway status.

The Route column uses these labels: Req. for a required specialist review; N/R when the review does not apply and the reason is accepted; and Reuse when current evidence has been accepted after checking its applicability. Stream. means the approved Tier 1–2 pathway, while Full requires the complete set of applicable evidence. Triage determines what the review must cover; Unclear means that scope is still undecided. A focused review addresses a particular change or issue while confirming that the other assumptions remain valid. Where appropriate, the reason field beside the required review describes this limited scope.

***Table L24**. Required reviews, reviewers, and routes.*

Review	When it is needed	Route	Reviewer and what the review covers Decision before which the review must be complete Reason or evidence link Owner and deadline for unresolved items
6.1 Data governance and privacy	Required for every use; detail should match data sensitivity and likely harm	☐ Req. ☐ Reuse	

Review	When it is needed	Route	Reviewer and what the review covers Decision before which the review must be complete Reason or evidence link Owner and deadline for unresolved items
6.2 Cybersecurity and architecture	Every use needs an initial cybersecurity review. A full architecture and security review is required for a new or changed device, software interface (API), cloud service, credential, integration, sensitive-data route, or automated action, or an unresolved concern.	☐ Triage ☐ Full ☐ Reuse ☐ Unclear	
7.1 Clinical safety and care quality	Care, clinical record, access, triage, diagnosis, treatment, or professional judgment	☐ Req. ☐ N/R ☐ Reuse ☐ Unclear	
7.2 Patients, equity, accessibility, Indigenous governance	Patients/public; access or service eligibility; employment decisions; personal communication; consent/choice; or Indigenous people, communities, or data	☐ Req. ☐ N/R ☐ Reuse ☐ Unclear	

Review	When it is needed	Route	Reviewer and what the review covers Decision before which the review must be complete Reason or evidence link Owner and deadline for unresolved items
7.3 Evidence and local assurance	Initial Tier 3 and Tier 4 approvals require full review. Only eligible Tier 1 and Tier 2 uses may use streamlined review through a qualified pathway. Any later focused review or reuse of evidence requires a recorded reason and confirmation that the other assumptions remain unchanged.	☐ Stream. ☐ Full ☐ Reuse ☐ Unclear	
7.4 Feature-specific review	System-capability checks apply to every enabled or requested function. Each check that does not apply requires an accepted reason; choosing a system type does not make every check relevant.	☐ Req. ☐ Reuse	
8 Operational readiness and lifecycle oversight	Required before operational use, restart, expansion, or a material workflow change. An earlier decision may identify the plan and the evidence still needed before release.	☐ Req. ☐ N/R ☐ Reuse ☐ Unclear	

Review	When it is needed	Route	Reviewer and what the review covers Decision before which the review must be complete Reason or evidence link Owner and deadline for unresolved items
9.1 Legal, regulatory, professional, and records	Required for Tier 3 or Tier 4, and for uses involving care, an official record, access, employment, rights, billing, research, or regulated activity. For other uses, required when it is unclear which rules apply.	☐ Req. ☐ N/R ☐ Reuse ☐ Unclear	
9.2 Procurement, vendor, and contract	Required for a new or changed purchase, partnership, provider, contract, or outside service on which the use depends. An earlier review may be reused only when its scope and version still match.	☐ Req. ☐ N/R ☐ Reuse ☐ Unclear	

Review	When it is needed	Route	Reviewer and what the review covers Decision before which the review must be complete Reason or evidence link Owner and deadline for unresolved items
5.2 Vendor POC authorization	Important unresolved questions about suitability, performance, integration, or workflow fit must be identified before a purchase or configuration commitment. A separate vendor POC is needed if adequate existing evidence or an equivalent authorized evaluation cannot answer them. If it is omitted, the procurement owner and responsible governance authority must record the evidence and their reasons under local policy.	☐ Req. ☐ N/R ☐ Unclear	Whether a POC is required and the supporting evidence; procurement owner; responsible governance authority; reason for the POC or approved omission; evaluation lead and specialist reviews; owner and due date for unresolved items:
External decision	Regulator, ethics board, shared network, insurer, or other authority	☐ Req. ☐ N/R ☐ Reuse ☐ Unclear	

Review	When it is needed	Route	Reviewer and what the review covers Decision before which the review must be complete Reason or evidence link Owner and deadline for unresolved items
Independent evidence review	Required when harm could be severe, the system is novel or autonomous, the developer controls most evidence, or local modification is substantial.	☐ Req. ☐ N/R ☐ Unclear	The record links to the independent review determination below and records when its conclusion is due.
5.3 Research-to-operations transition	Required when research output, evidence, software, data access, or workflow becomes part of routine service.	☐ Req. ☐ N/R ☐ Reuse ☐ Unclear	Reviewer; ethics review or other reference; owner and due date:

***Table L25**. Activities permitted while review continues.*

Work allowed during review Anything beyond planning or contract discussion needs a recorded authorizing decision, relevant clearances, limits, and expiry. Submitting this form or arranging reviews does not grant that permission.	☐ None ☐ Planning or contract discussion ☐ Research ☐ Retrospective review of existing data ☐ Simulation ☐ Silent evaluation—outputs do not guide care or operations ☐ Current authorized use
Limits during review The record describes the permitted activity and how outputs will be prevented from guiding care or operations where required. It includes data and system access, owner, dates, incident reporting procedure, and expiry.	

***Table L26**. Screening and review-route sign-off.*

Role	Name	Signature or electronic confirmation	Date
Governance lead			

4.3 Independent review requirements

***Table L27**. Independent review determination.*

Independent review field	Response or evidence reference
Is independent review required?	☐ Required ☐ Not required Reason:
Trigger and reviewer The record states what triggers independent review: severity, novelty, autonomy, developer-controlled evidence, or substantial local modification. It names the reviewer and the decision that needs their conclusion.	
Independent review: evidence and conclusion The record states what the reviewer examined, explains their independence, and discloses conflicts. It includes the conclusion, evidence link, and date. Each unresolved question has a named owner and an identified decision that must wait. A pending review is not a favourable finding.	

5 Sourcing and service support

The requesting team, internal technical owner, and operational owner assess how the hospital will provide and support the service. Their comparison of realistic options covers the time required, lifetime cost, and control the hospital would retain. It establishes whether the hospital can operate and oversee the service, maintain its configuration group within approved limits, and leave the arrangement safely if needed. Procurement participates when buying, partnering, or combining routes. The choice requires supporting local evaluation, enforceable responsibilities, and funded capacity.

Each material component has an assigned route: Build when the hospital develops or modifies it, Buy when a supplier provides it, including as a managed service, and Partner when design or decision authority is shared. A service may combine routes. The component map names who performs each task and which hospital owner checks and accepts the work. The hospital remains accountable for local use and decisions. It records the evidence due, response time, funding, backup arrangements, and what the hospital will do if the responsible party fails to deliver. Every duty needs a capable owner before the decision that depends on it. The common checks apply once, followed by the checks for each applicable route. **Table L28** records the route; **Table L29** records the reasons for choosing it.

***Table L28**. Sourcing route decision.*

Selected route Undecided is acceptable only during early planning with no purchase commitment or live-data use.	☐ Build ☐ Buy/configure, including a managed service ☐ Partner/co-develop ☐ Mixed components—component map required ☐ Undecided Details and confirmed component mapping:

***Table L29**. Rationale for the sourcing choice.*

Reason for the sourcing choice The assessment compares realistic options for capability, time, cost over the service's life, local control, evidence access, integration, monitoring, and exit. The assessment draws on the recorded comparison of the AI proposal with current practice and non-AI alternatives.	

Each material component or safeguard has a separate row in **Table L30**, with its name or identifier and Build, Buy, or Partner route in the first column. Separate rows are needed whenever the person doing the work, hospital owner, evidence, deadline, safe alternative, or funding differs. A service that combines routes requires the applicable checks for each, linked to the component names or identifiers. The label "mixed" cannot replace an assignment of responsibility. **Table L31** records the common sourcing, control, and support requirements.

***Table L30**. Component responsibilities and evidence.*

Component or safeguard / Build, Buy, or Partner	**Who does the work and controls this component**	**Hospital owner / decisions the hospital retains**	**Evidence required at this stage / deadline / safe alternative / funding**
Base model or application, its origin and development history			
Data processing and flow, inputs, required formats (schemas), and reference range			
Prompt, template, safety restriction (guardrail), retrieved reference content, or threshold			
System connections, identity controls, EHR write-back, tools, and permissions			
Release and change approval; fixed version or release identifier			
Local evaluation, reference cases, and tests of whether previously accepted functions still work (regression tests)			
Local use, required user practice, workflow, and support			
Performance and quality monitoring			
Service availability and defect correction			

Component or safeguard / Build, Buy, or Partner	Who does the work and controls this component	Hospital owner / decisions the hospital retains	Evidence required at this stage / deadline / safe alternative / funding
Incident investigation, correction, and prevention of recurrence			
Rollback or safe disabling and recovery			
Maintaining service safely during disruption			
Data and record export			
Exit and retirement			

***Table L31**. Common sourcing, control, and support requirements.*

Value and costs over the service's life The recorded value assessment identifies the implementation, operating, monitoring, and exit costs for this sourcing option over the intended period of use or contract term.	
Hospital control of the service The record links to the exercise, test, or enforceable term showing that the hospital can pause or disable use, inspect logs and versions, investigate problems, and correct affected records or actions.	

Leaving the arrangement safely Evidence demonstrates that the hospital can roll back or safely disable the service, export required records, revoke access, and verify deletion while maintaining a safe alternative.	
Hospital decision-makers The record identifies who may decide the intended use, tier, acceptance of remaining risk, release, expansion, pause, restart, and exit. Accountable owners are distinguished from people with formal approval powers.	
Assigned work The record links to the component map naming who performs each task and the hospital owner who checks it. It identifies any unassigned duty and the decision that must wait until it has an owner.	
Configuration group and approved operating limits The record names the group, lists its permitted variants, and links to the comparison showing that they are equivalent for this use. It explains which changes fit within the approved operating limits and which need review.	

Number of configuration groups The record provides the locally approved maximum and the current and proposed number of groups. The framework suggests one group per service for capacity planning; this is not a measured technical limit. An increase requires a reason based on capacity and evidence that support is funded. It includes the named authority's approval and a review or expiry date.	
Staffing and funding for review and operation The record includes the commitments for evidence review, monitoring, support, incident response, and rollback or safe disabling at the proposed number of groups and scale of use.	
When the group limit is reached The record states whether the authorized response is consolidation, retirement, declining the proposed variation, or an exception supported by additional funded capacity. It explains how an exception preserves core safeguards and mandatory duties.	
When more capacity is needed The record defines the workload, backlog, or service-level threshold that requires more staff or resources, a process change, or automation that supports the work. It identifies who decides and the response deadline.	

Current capacity assessment The record links to a dated comparison of workload, backlog, timeliness, errors, and overdue actions against that threshold. It includes the response required, if any.	
Tracking and automating governance work The record explains how the intake and review process tracks this use, which work is automated, and which decisions require an authorized reviewer.	
Controls on governance automation The record identifies the workflow or software version, its owner, and where unresolved or exceptional cases are queued for review. It links to the audit trail of actions, changes, and decisions. It states which decisions an authorized person must make.	
Checking governance automation The record includes the approved sampling plan, threshold for material errors, and rule for suspending automation. It describes the manual process used when it is unavailable or suspended.	

Responsibilities for change and drift For each relevant change or drift, the record identifies who is responsible for detecting, reviewing, investigating, and correcting it, and who funds the response. Drift means a departure from the conditions that supported approval. The record identifies any independent review required.	
Supplier evidence and deadlines The record states what evidence the supplier or partner must provide, who receives it, and the deadlines for responding and correcting problems.	
If the responsible party does not resolve a problem The record explains when the hospital will escalate, restrict use, roll back, or exit if the responsible party fails to detect, explain, or correct the problem.	
Testing controls for change and drift The record links to the test record for each safeguard. This includes the component, owner, test method, pass criterion agreed in advance, latest result, date, and frequency of testing.	
Supplier or partner escalation where applicable The record explains how an internal investigation leads to a supplier response and governance escalation. It states the response deadline and when executive leaders or the board must become involved.	

5.1 Checks by sourcing route

The component map determines which Build, Buy, and Partner tables apply. The shared checks apply once. A service may require more than one route table: for example, a purchased model connected through hospital-built software needs Buy checks for the model and Build checks for the software. Common evidence is linked rather than copied. The Met, Open, and N/A rules at the start of this form apply; evidence is due before the decision that depends on it, and a reviewer must accept any N/A reason. Differences in evidence or status require a separate row identifying the component. Each Open item has a Pending, Failed, Unknown, or Expired status and a linked action record identifying its owner, due date, and the decision that must wait.

5.1.1 Shared checks

The component map and common checks above are the single record of assigned duties, funding, local control, continuity, and exit. The two checks in **Table L32** apply once to the whole service. The route tables refer back to that record. The sourcing decision owner records the review conclusion below and signs **Table L36**.

***Table L32**. Shared checks for component routes and responsibilities.*

Check and required evidence	Result and evidence reference
S1 Component routes The component map identifies every material component by name or identifier and assigns a Build, Buy, or Partner route. For a mixed service, it identifies every route used and links each component to the applicable route checks.	☐ Met ☐ Open ☐ N/A Evidence and reviewer/date: Open category and action reference: N/A reason and acceptance:
S2 Component responsibilities Each component has an accountable hospital owner and assigned responsibility for evidence, monitoring, correction, rollback or safe disabling, maintaining service during disruption, and exit. The completed component map provides the single record of these assignments.	☐ Met ☐ Open ☐ N/A Evidence and reviewer/date: Open category and action reference: N/A reason and acceptance:

5.1.2 Build route

Table L33 applies only to components marked Build. Component names or identifiers from the component map: _____. Local development or modification carries the corresponding Build duties, even when the original product was purchased.

***Table L33**. Assurance for hospital-built components.*

Check and required evidence	Result and evidence reference
B1 Development and maintenance capacity available The common staffing and funding record confirms the required roles, backups, working environments, and sustained funding.	☐ Met ☐ Open ☐ N/A Evidence and reviewer/date: Open category and action reference: N/A reason and acceptance:

Check and required evidence	Result and evidence reference
B2a Release and data history recorded The current record shows how the release was produced and how data were obtained and processed (release and data lineage).	☐ Met ☐ Open ☐ N/A Evidence and reviewer/date: Open category and action reference: N/A reason and acceptance:
B2b Release tests current Current regression tests show that accepted functions still work after changes. Tests also confirm the required structure, fields, and format of incoming and outgoing data.	☐ Met ☐ Open ☐ N/A Evidence and reviewer/date: Open category and action reference: N/A reason and acceptance:
B3a Operation and continuity ready Evidence confirms that arrangements are ready for change approval, monitoring, rollback or safe disabling, and maintaining service safely during disruption. The record links to the shared sourcing and support records.	☐ Met ☐ Open ☐ N/A Evidence and reviewer/date: Open category and action reference: N/A reason and acceptance:
B3b Maintenance handover ready A named successor team or person can take over maintenance. The record links to evidence that the handover arrangements are ready.	☐ Met ☐ Open ☐ N/A Evidence and reviewer/date: Open category and action reference: N/A reason and acceptance:

5.1.3 Buy route

Table L34 applies only to components marked Buy. Component names or identifiers from the component map: _____. A managed service belongs here unless design or decision authority is shared.

***Table L34**. Assurance for purchased products and managed services.*

Check and required evidence	Result and evidence reference
Y1a Supplier evidence current The record links to current vendor evidence and contacts.	☐ Met ☐ Open ☐ N/A Evidence and reviewer/date: Open category and action reference: N/A reason and acceptance:

Check and required evidence	Result and evidence reference
Y1b Hospital control demonstrated The record links to evidence that the hospital controls its local configuration and monitoring, including when the vendor hosts or operates the service.	☐ Met ☐ Open ☐ N/A Evidence and reviewer/date: Open category and action reference: N/A reason and acceptance:
Y2 Supplier operating duties The component map identifies who handles updates, defects, incidents, corrective action, and support, with links to the terms assigning those duties. It provides the single responsibility list.	☐ Met ☐ Open ☐ N/A Evidence and reviewer/date: Open category and action reference: N/A reason and acceptance:
Y3a Service controls usable Tests demonstrate rollback or safe shutdown, as appropriate to the service, and effective control of enabled or disabled features. Linked common control records and terms preserve those capabilities.	☐ Met ☐ Open ☐ N/A Evidence and reviewer/date: Open category and action reference: N/A reason and acceptance:
Y3b Remedies and exit rights usable Terms and test results show that remedies for unmet obligations and exit rights can be used in practice, with links to the common control and exit records.	☐ Met ☐ Open ☐ N/A Evidence and reviewer/date: Open category and action reference: N/A reason and acceptance:

5.1.4 Partner route

Table L35 applies only to components marked Partner. Component names or identifiers from the component map: _____. The record states the shared design or decision authority. Calling a supplier a partner does not establish shared responsibility.

***Table L35**. Assurance for jointly developed or governed components.*

Check and required evidence	Result and evidence reference
P1 Decision and change authority Each shared component has a linked record showing who may make each decision, how changes are reviewed jointly, and the hospital's authority to release the service locally.	☐ Met ☐ Open ☐ N/A Evidence and reviewer/date: Open category and action reference: N/A reason and acceptance:

Check and required evidence	Result and evidence reference
P2 Evidence, disputes, and rights The record links to the agreed arrangements for access to evidence, escalation of disputes, and intellectual-property, publication, and data rights.	☐ Met ☐ Open ☐ N/A Evidence and reviewer/date: Open category and action reference: N/A reason and acceptance:
P3 When the partnership ends The record links to the continuity and transition plan for the end of the partnership. The plan shows who carries out the work, how it is funded, and why the arrangements are workable.	☐ Met ☐ Open ☐ N/A Evidence and reviewer/date: Open category and action reference: N/A reason and acceptance:

Supporting file or link:

Review conclusion: ☐ Ready ☐ Ready with required actions ☐ Not ready

***Table L36**. Sourcing review sign-off.*

Sourcing decision owner name and role	Signature / electronic confirmation	Date

5.2 Vendor evaluation

A separate, time-limited vendor POC is required when important questions about suitability, performance, integration, or fit with hospital work need answers before a purchase or configuration commitment. It may be omitted if adequate existing evidence or an equivalent authorized evaluation answers those questions. The procurement owner and responsible governance authority must document the evidence and their reasons in line with local policy. An internal build with no vendor selection does not need a vendor-selection POC, but every use still needs local assurance. **Section 5.2** records the procurement question, products or configurations being compared, evaluation setting, dates, acceptance criteria, and stop rules. POC results must not guide care or operations. A live pilot needs approval for limited operational use and production confirmation. Access to production data or live systems requires separate linked clearances before evaluation starts.

Table L37 records the evaluation scenarios, data, evaluators, support, access limits, and closure plan. These details help later reviewers understand what the results can and cannot establish. The applicable specialist clearances and safeguards for data, access, safety, cybersecurity, incidents, and stopping must be complete before evaluation starts. Conditional approval may leave later deliverables outstanding, but cannot defer these prerequisites. POC evidence may support procurement or local-assurance review when it remains applicable. **Table L38** records authorization to evaluate and its limits.

***Table L37**. Vendor POC scope and prerequisite checks.*

Check and information to provide	Status and evidence reference
Procurement question and POC route resolved The record includes the purchase or configuration decision to be made, the options, and the questions to be examined in the POC. It states which results would stop or change procurement. Omitting the POC requires adequate existing evidence or an equivalent authorized evaluation to support the reasons. The procurement owner and responsible governance authority must accept those reasons under local policy.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Evaluated product and configuration The record includes the vendor, product, model, release, features, hosting, and integration.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Evaluation access defined The record specifies the test environment, permitted data access, credential owners and permissions, connection points (endpoints), and separation from other systems. It links to the controlled credential record; passwords and secrets must not appear in this form.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Evaluation schedule The record includes the evaluators, planned sessions, authorization start and end dates, and expiry.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Evaluation support and cost The record includes vendor and internal support, response times, staff time, estimated cost, and funding owner.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Fair comparison The assessment compares vendors against the same need, scenarios, comparator, and measures. The record explains unavoidable differences between products and states which evaluation conditions must stay the same.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:

Check and information to provide	Status and evidence reference
Evaluation data Evaluation uses synthetic or appropriately de-identified data where possible. Personal health information or live-system access requires a recorded justification, controls, expiry, and linked separate authorization.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Vendor access and data handling The record includes vendor access, other companies it uses to process data (subprocessors), hosting location, uses beyond the evaluation purpose, and any further transfers.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Plan to end the evaluation ready before approval Evidence confirms that the plan covers removal of temporary access and return, retention, or deletion of data and records. It must state the deadlines, name the responsible person and verifier, and identify the evidence needed to confirm completion.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Prerequisite reviews The record links to the privacy, cybersecurity, safety, legal, data-governance, and research decisions required for the proposed data, access, and connection points.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Evaluation kept separate from operations Verification confirms the required separation from operational systems and that access is limited to the authorized evaluation. POC outputs cannot guide care or operations.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Evaluation activity traceable Evidence confirms that activity is logged and incidents can be reported. The record links to evidence that both arrangements work.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:

Check and information to provide	Status and evidence reference
Acceptance criteria set in advance Each criterion has a recorded measure and unit, sample or subgroup, threshold, owner, approval date, and evidence reference. The criteria, including thresholds, were approved before evaluation began.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Stop and change rules The record specifies which conditions require immediate suspension and which changes require reauthorization before evaluation can resume.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Roles and authority The record names the procurement owner, evaluation lead, evaluators, specialist reviewers, and decision-maker. It discloses any conflicts.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:

5.2.1 Vendor POC authorization

Table L38. *Vendor POC authorization.*

Decision field	Decision or record
Vendor POC decision	☐ Authorize the proposed activity and limits ☐ Authorize with later conditions ☐ Return for revision ☐ Hold ☐ Reject
Authorization record The record includes the authorized activity, agreement, cost, and safeguards that must already be in place. It also includes any later conditions and their owners, the decision authority, date, and expiry. It links to separate data or access clearances. This permits the evaluation only, not a production purchase or operational use.	

5.2.2 Conduct during the POC

The POC authorization covers only the recorded scope and conditions. A material change requires approval before evaluation resumes. The session record identifies the evaluator, product and configuration, scenario, data, and result. This must allow reviewers to trace errors, deviations, vendor interventions, and changes. Reaching a stop threshold or

identifying a suspected incident requires the affected work to stop and the incident procedure to begin. **Table L39** tracks POC conduct, **Table L40** records results and close-out, and **Table L41** records the procurement decision.

***Table L39**. Vendor POC conduct and traceability.*

Check and information to provide	Status and evidence reference
Sessions and scenarios The record includes each session, evaluator, product and configuration, scenario, data, and result.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Deviations and amendments The record includes errors, deviations, amendments, and vendor intervention.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Stops and incidents The record includes stop events, suspected incidents, and the action taken.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Evaluator independence The record identifies who selected the cases, controlled evaluation data and settings, verified the results, and made the procurement decision. It discloses vendor involvement and conflicts that could affect independence.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:

5.2.3 POC results and closure

***Table L40**. Vendor POC results and close-out.*

Evaluation finding to record	Response or evidence reference
Evaluation completed The record includes the product, release, configuration, environment, data, evaluators, dates, and sessions actually completed.	
Actual cost and support used The record includes staff time, vendor and internal support, and actual cost.	

Evaluation finding to record	Response or evidence reference
Deviations, stops, and unresolved issues The record includes every deviation, stop, incident, and unresolved item.	
Completeness check The assessment compares what was planned with the sessions, configurations, and scenarios actually completed. The record explains deviations, changes, stops, incidents, and missing evidence.	
Independent verification The record identifies who verified completeness, independence, and the final evidence set. It includes the date and conclusion.	
Criteria and results For every criterion agreed before evaluation, the assessment compares the required result with the observed result. This includes relevant subgroup and difficult-case findings. The record includes Pass, Fail, or Insufficient, with the uncertainty and evidence reference. It explains any criterion that does not apply.	
Usability and user burden The record includes important errors, how readily they could be detected, the work needed to use the system, and the burden observed.	
Support performance and limitations The record includes support response, unresolved issues, and service limitations.	
System connection and data-exchange results The record includes the results for integrations, data flows, and logging.	
Investigation and correction capability The record states whether the hospital could obtain evidence, investigate an issue, and verify correction.	
Rollback and service control The record states whether the evaluated configuration could be disabled or rolled back, by whom, and within what time.	

Evaluation finding to record	Response or evidence reference
Data export, deletion, and exit The record includes findings on data export, verified deletion, and ending the service arrangement.	
Differences from the proposed live service The assessment compares the tested configuration with the proposed live service. The comparison covers contract, hosting, integration, permissions, monitoring, support, users, and scale. The record states what additional evidence is needed because of those differences.	
Cost and funding The assessment compares acquisition, integration, operation, monitoring, support, change, continuity, and exit costs, identifying the funding source and term.	
Contract suitability The record explains how important contract gaps, dependencies, remedies for unmet obligations, and negotiation conditions affect the purchase decision.	
Removal of temporary access The record documents removal of temporary accounts, credentials, connection points (endpoints), integrations, and vendor access. It identifies who verified completion and when.	
Data and records at the end of evaluation The record documents the approved return, retention, or deletion of data, copies, logs, and outputs, while preserving the evaluation record.	
Issues to carry into the next review The record lists contract, cost, incident, evidence, or production issues that must be addressed in contracting or local assurance. Every unresolved item has an owner and deadline.	

5.2.4 Procurement decision

Table L41. *Procurement decision after the POC.*

Decision field	Decision or record
Procurement decision	☐ Stop ☐ Seek information ☐ Change specification ☐ Compare another option ☐ Negotiate conditions ☐ Continue procurement
Decision record The record explains how the results against the criteria and the remaining uncertainty support the procurement decision. It includes unresolved conditions and their owners, the decision authority, date, and any change to the component map. Continuing procurement does not extend POC access or permit operational use.	

5.3 Research into routine use

Table L42 is required when research software, data access, evidence, or a workflow is proposed for routine service. Research approval does not authorize operational use. Reviewers may reuse evidence that remains relevant, but each activity still needs its own authorization. The transition record identifies the evidence that can carry forward, what has changed, and what is still needed for routine use. Decision stage and service status remain separate entries, as explained at the start of this form. All required operational clearances must be in effect before live use begins.

Table L42. *Research-to-operations transition checks.*

Check and information to provide	Status and evidence reference
Research and ethics authority The record includes the ethics reference and status and any required consent, patient-information, or Indigenous/community authority.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Limits while the use remains research only The record states how research outputs are prevented from guiding routine care or operations before operational authorization.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:

Check and information to provide	Status and evidence reference
Research incidents and deviations The record links to any incident or deviation relevant to the proposed transition. It includes what was decided and the action taken or still required.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Proposed live service defined Evidence confirms that the linked record defines the proposed configuration, use, and limits. It must cover data access, integrations, authority to take actions, credentials, user instructions, support, monitoring, safe alternatives, and stopping.	☐ Met ☐ Open Evidence reference or finding:
Research evidence assessed for live use The record explains why the research evidence applies to the proposed live service. It identifies important differences, remaining gaps, and any new acceptance criteria.	☐ Met ☐ Open Evidence reference or finding:
Purpose in operational use The record states the routine-service purpose and the decision or action it supports.	☐ Met ☐ Open Evidence reference or finding:
Data-use authority The record includes the authority to use data for the new purpose and any decision about notice, consent, or choice.	☐ Met ☐ Open Evidence reference or finding:
Regulatory and professional status The record includes the applicable regulatory and professional requirements. It states any status that must be established to meet them. It identifies the reviewer, conclusion, and supporting evidence.	☐ Met ☐ Open Evidence reference or finding:
Official-record duties The record identifies any official record the service will create or change and the duties for retaining and correcting it. If none is affected, the record includes the reviewer's accepted conclusion.	☐ Met ☐ Open Evidence reference or finding:

Check and information to provide	Status and evidence reference
Privacy conclusion The record links to the required privacy review and its conclusion.	☐ Met ☐ Open Evidence reference or finding:
Cybersecurity conclusion The record links to the required cybersecurity review and its conclusion.	☐ Met ☐ Open Evidence reference or finding:
Clinical-safety conclusion The record links to the required clinical-safety review and its conclusion.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Equity conclusion The record links to the required equity review and its conclusion.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Accessibility conclusion The record links to the required accessibility review and its conclusion.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Legal conclusion The record links to the required legal review and its conclusion.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Research ownership and publication rights The record includes ownership of data, models, outputs, and intellectual property, and the applicable publication and review rights.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:

Check and information to provide	Status and evidence reference
Funding and conflicts The record discloses funding, in-kind support, investigator or institutional interests, and the decision on how any conflicts will be managed.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Supplier or partner duties during live use The record includes obligations for support, change notices, incidents, audits, corrective action, service continuity, and exit.	☐ Met ☐ Open ☐ N/A Evidence reference or finding:
Owner acceptance of production duties The record includes each owner's dated acceptance of their production duties: the clinical/end-user owner, internal technical owner, and operational owner.	☐ Met ☐ Open Evidence reference or finding:
Registry and dossier update The record identifies who updated the registry and dossier, when, and for which decision. The operational status is updated after governance approval and again after production confirmation.	☐ Met ☐ Open Evidence reference or finding:

6 Data protection and system design

6.1 Data governance and privacy

The data governance and privacy reviewer checks every use. Review depth must match data sensitivity and possible harm. The review follows the data from its source through every place it is processed, stored, or sent, including later uses. It compares the data inventory and actual flows with approved purposes and permissions. Each difference, exception, or missing safeguard is linked to the decision it affects. A plan may be sufficient for the current activity, but safeguards must be verified as working before the hospital grants the access or release that depends on them. **Table L43** records the review, and the reviewer signs **Table L44** to confirm the conclusion.

***Table L43**. Data governance and privacy review.*

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Approved data inventory The record includes every data source and element, purpose, recipient, transfer, storage location, supplier, and downstream system.	☐ Met ☐ Open ☐ N/A	
Check of the actual data flow An independent reviewer must compare the observed data flow with the approved map in the authorized test or release environment. The record identifies that environment and any checks still needed in production.	☐ Met ☐ Open ☐ N/A	
Differences in data flow Each difference has a recorded decision, owner, and supporting evidence. The record explains how it will be addressed. It includes temporary safeguards and the decision that must wait until the difference is resolved.	☐ Met ☐ Open ☐ N/A	
Basis for each data use The record identifies the legal, ethical, consent, or contractual basis for collection, access, linkage, recording, disclosure, retention, and any use beyond the approved purpose.	☐ Met ☐ Open ☐ N/A	
Input data quality Each input whose failure could affect a decision has a recorded source and purpose, quality measure, baseline, acceptance threshold, checking frequency, and data owner.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
When input data fail a quality check The record provides the latest result and how errors will be corrected. It states the action or stop rule that applies when the threshold is not met.	☐ Met ☐ Open ☐ N/A	
External processing and later data use controlled The record identifies and sets limits on hosting, the provider's outside processing companies, jurisdictions, training or product-improvement use, further transfers, and sensitive-data settings.	☐ Met ☐ Open ☐ N/A	
Access to approved data and systems The record identifies the roles and current approvals. Evidence demonstrates that access is limited to what each role needs (least privilege); the current access-review results are linked.	☐ Met ☐ Open ☐ N/A	
Retention and legal holds The record links to the approved retention periods, storage locations, and legal holds. A legal hold requires information to be preserved beyond its usual deletion date.	☐ Met ☐ Open ☐ N/A	
Deletion or return The record includes the methods for deleting or returning data, the deadlines, the verifier, and the supporting evidence.	☐ Met ☐ Open ☐ N/A	
Exceptions to data controls Each exception has a recorded authorizer, expiry, temporary safeguard, and date for repeat testing.	☐ Met ☐ Open ☐ N/A	
Notice or consent for data use The record includes the basis of authority for the data use. It states who must receive information, what it must cover, and when it must be provided. It includes the effective date.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Providing required privacy information Delivery records confirm that the approved privacy information was provided in the required language and format.	☐ Met ☐ Open ☐ N/A	
If a person refuses data use Where required, evidence demonstrates that the approved refusal or opt-out process and equivalent non-AI route work in practice. Otherwise, the record links to the applicability decision and explains what people are told about their options.	☐ Met ☐ Open ☐ N/A	
Data correction route The record describes how an error can be corrected, who is responsible, and the required service level, including response time.	☐ Met ☐ Open ☐ N/A	
Initial privacy review The record links to the current privacy screening (triage). It states what it covers, its conclusion, the reviewer, and the date.	☐ Met ☐ Open ☐ N/A	
Full privacy or data-impact review, if required The record links to the privacy impact assessment or equivalent review. This includes approval, unfinished actions, and the next review or expiry date.	☐ Met ☐ Open ☐ N/A	
Minimum necessary data The record explains why each data element is needed. It shows how unnecessary inputs, outputs, logs, and derived data are excluded or reduced.	☐ Met ☐ Open ☐ N/A	
Ongoing data-flow checks established The record includes who compares actual flows with the approved map, how, and how often. It includes the latest result, exceptions, and response. This includes supplier confirmation or independent tests where relevant.	☐ Met ☐ Open ☐ N/A	

The affected data access or release is Not ready if a data route or permission is unknown, sensitive data could leave approved limits, or use beyond the original purpose remains unresolved. The same applies if access, retention, or

deletion cannot be controlled. The finding states whether any earlier preparation may continue; this does not permit the blocked access or release.

Supporting file or link:

Review conclusion: ☐ Ready ☐ Ready with required actions ☐ Not ready

***Table L44**. Data governance and privacy sign-off.*

Reviewer name and role	Signature / electronic confirmation	Date

6.2 Cybersecurity and technical design

The cybersecurity and technical reviewer assesses whether the service can be secured and monitored, stopped when necessary, investigated, and safely restored. The component inventory, threat assessment, and tests must cover the configuration under review, including user and software accounts, external services or components, and data routes. Each test record states where the test was performed, any differences from the intended setup, and what failed. Reliance on a test for release requires verification that the live setup matches the tested setup. **Table L45** records the review, and the reviewer signs **Table L46** to confirm the conclusion.

***Table L45**. Cybersecurity and technical design review.*

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Approved component inventory Evidence confirms that cloud services, software interfaces (APIs), devices, models, providers, and development, test, or live environments match the approved technical design.	☐ Met ☐ Open ☐ N/A	
Connections and data routes tested Evidence confirms that external connections and data routes have been tested in the authorized environment. The record links to the results.	☐ Met ☐ Open ☐ N/A	
Production matches the tested configuration Before live release, verification confirms that production connections and data routes match the tested configuration.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Differences from the approved design The record describes each difference, the decision on how to address it, the responsible owner, the temporary safeguard, and the evidence that it has been resolved.	☐ Met ☐ Open ☐ N/A	
Separation of development, testing, and production Evidence confirms that these environments are separated as required. This includes restrictions on live data and the controlled process for moving software into production.	☐ Met ☐ Open ☐ N/A	
Human access controls The record describes authentication, approval, access limited to what the role needs, and permitted actions. This includes access review, removal, and any time-limited exception.	☐ Met ☐ Open ☐ N/A	
Software accounts and secrets controlled Each service account and secret, such as an API key, has a confirmed owner and minimum permissions. Arrangements for secure storage, replacement, expiry, and revocation are verified. Secrets must not appear in this form.	☐ Met ☐ Open ☐ N/A	
Tests of unauthorized access The record describes tests of attempts to gain access or take actions without permission. It provides the latest results, failed controls, verifier, and date.	☐ Met ☐ Open ☐ N/A	
Access exceptions Each exception specifies which access it permits and its limits. The record includes the alternative safeguard, approving authority, expiry, owner, and date for repeat testing.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Threat assessment matches the configuration The record links to the dated assessment of security threats and possible harm to the service or its data (the threat model). Evidence confirms that it matches the configuration under review. The assessment is checked again against the intended release.	☐ Met ☐ Open ☐ N/A	
Relevant attacks tested Tests cover data leakage and attempts to redirect the system through input or retrieved instructions (prompt injection), malicious or deliberately corrupted inputs and reference content, misuse of privileges or permitted scope, unsafe use of connected tools, and service compromise.	☐ Met ☐ Open ☐ N/A	
Cybersecurity finding recorded The record includes the specialist's finding, unresolved conditions, and the decisions that depend on them.	☐ Met ☐ Open ☐ N/A	
Remaining cybersecurity risk accepted Evidence confirms that the authorized decision-maker has accepted the remaining risk within their powers. The record links to that decision and its conditions.	☐ Met ☐ Open ☐ N/A	
Correcting failed security tests Every failure has a recorded corrective action, owner, due date, and result of repeat testing.	☐ Met ☐ Open ☐ N/A	
External dependencies The record identifies providers, models, libraries, outside companies used by providers, hosting, support periods, known vulnerabilities, and how updates or changes are made.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Logs for security and incident review Evidence demonstrates that the hospital can use the logs to identify the user and version, important inputs, outputs, and actions, privileged changes, alerts, and the incident timeline for the required period.	☐ Met ☐ Open ☐ N/A	
Stop or restriction control tested Tests show that authorized users can stop or restrict the affected function within the required time.	☐ Met ☐ Open ☐ N/A	
Stop control verified in the release environment Before operational access, verification confirms that the tested stop or restriction control works in the environment intended for live use.	☐ Met ☐ Open ☐ N/A	
Capacity during downtime Evidence demonstrates that the manual or alternative work process has enough capacity for the approved period of downtime.	☐ Met ☐ Open ☐ N/A	
Backup and restore test The record links to the backup and restoration tests and their results, including whether data loss stayed within the permitted limits.	☐ Met ☐ Open ☐ N/A	
Rollback or safe shutdown works Tests confirm a return to an approved working version, or safe disabling with a functioning alternative. The record includes any technical limitation and the accepted continuity arrangement.	☐ Met ☐ Open ☐ N/A	
Encryption tests The record identifies the data routes and environments tested. It explains how encryption was tested while data were being sent and while stored. It provides the results, verifier, and date.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Encryption-key controls The record explains the tests of who holds and can access encryption keys, how keys are replaced, and how they can be recovered. It provides the results, verifier, and date.	☐ Met ☐ Open ☐ N/A	
Encryption or key-control exceptions The record identifies who approved the exception, its expiry, the person responsible for correction, and their deadline. It includes the result of repeat testing.	☐ Met ☐ Open ☐ N/A	

The service remains Not ready for release while a critical security issue is unresolved or a sensitive-data route is unknown. A Tier 3 or Tier 4 service must also remain Not ready if it lacks a tested way to stop and recover through an approved version or safe alternative.

Supporting file or link:

Review conclusion: ☐ Ready ☐ Ready with required actions ☐ Not ready

***Table L46**. Cybersecurity and technical design sign-off.*

Reviewer name and role	Signature / electronic confirmation	Date

7 Safety, people, and evidence

7.1 Clinical safety and care quality

The clinical safety and quality reviewer examines any use that may affect care, clinical records, patient access, or professional decisions, using the agreed description in **Sections 3 and 4**. The review traces each serious hazard through the actual workflow. Testing must establish whether the safeguards and human review can prevent or contain its consequences under realistic conditions. Readiness is judged against safety criteria approved before evaluation begins, with tests of the safe alternative and escalation procedure included. **Table L47** records the review, and the reviewer signs **Table L48** to confirm the conclusion.

***Table L47**. Clinical safety and care quality review.*

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Clinical purpose and limits The clinical decision or workflow, population, setting, version or setup, users, and requested governance decision match the agreed intended use and review scope.	☐ Met ☐ Open ☐ N/A	
Serious hazards and safeguards assessed Each hazard has an explanation of how it could arise and affect people. The record includes the possible harm, safeguard, remaining risk, responsible owner, and monitoring or stop signal. It links to evidence that each safeguard works.	☐ Met ☐ Open ☐ N/A	
Human review is feasible Evidence confirms that reviewers have the source information, expertise, and time needed to assess the output, the authority to disagree with it, and a procedure for escalating concerns.	☐ Met ☐ Open ☐ N/A	
Review and override records tested Test results show that reviews and decisions to override AI output are recorded as required.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Clinical advice and uncertainty are clear Advice or a differential diagnosis presents alternatives, reasons or sources, and limitations clearly. The record defines when the system should refuse a request or withhold output, and what safe action should follow.	☐ Met ☐ Open ☐ N/A	
Testing with frontline staff Testing with representative roles and shifts assesses disagreement, missing or wrong output, downtime, and handover against the expected response and pass criteria.	☐ Met ☐ Open ☐ N/A	
Results of testing with frontline staff The record includes the participants, observed result, unresolved issue, responsible owner, and retest date.	☐ Met ☐ Open ☐ N/A	
Clinical-safety criteria agreed in advance The record includes the safety measures, acceptance criteria, warning levels, and mandatory stop thresholds that were approved before reviewers saw the results.	☐ Met ☐ Open	
Test of the safe alternative workflow The record includes the alternative, test scenario, result, responsible owner, and date.	☐ Met ☐ Open	
Test of patient-safety escalation The record includes the exercise, result, unresolved action, responsible owner, and retest date.	☐ Met ☐ Open	

The affected activity is Not ready if a serious hazard is not controlled, required human review cannot work as intended, a safety criterion fails, or the safe alternative or escalation procedure cannot be used in practice.

Supporting file or link:

Review conclusion: ☐ Ready ☐ Ready with required actions ☐ Not ready

***Table L48**. Clinical safety and care quality sign-off.*

Reviewer name and role	Signature / electronic confirmation	Date

7.2 Patients, equity, accessibility, and Indigenous governance

In **Table L49**, patient, equity, and accessibility reviewers assess whether people can understand and use the service, participate in decisions that affect them, and obtain review or correction. The review establishes whether information reaches people and is understood, and whether required choice and complaint processes work. It examines differences in access or outcomes between locally relevant groups and records how participation influenced the proposal. Where Indigenous people, communities, or data may be affected, the applicable Indigenous-led process establishes requirements and decision authority. The record includes every required protocol, agreement, or approval.

***Table L49**. Patient, equity, accessibility, and Indigenous governance checks.*

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Service information provided Information is provided in approved languages and formats. It explains what the AI does, its purpose and limits, who remains responsible, data use, available choices, and how to seek correction or complain.	☐ Met ☐ Open ☐ N/A	
Understanding and accessibility tested The record includes the intended audience, delivery timing, and applicable accessibility standard. It provides the test method and participants, acceptance threshold, result, owner, and update triggers.	☐ Met ☐ Open ☐ N/A	
Notice, consent, and choice requirements The record states each requirement and records its authorized decision-maker, supporting reference, effective date, expiry, and any outstanding action.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Delivering approved information Evidence confirms that the information reaches people at the right time, in the required language and accessible format, and that delivery is documented.	☐ Met ☐ Open ☐ N/A	
Refusal, opt-out, and alternatives Testing covers any required refusal or opt-out process. A practical non-AI route is available and tested where feasible or required. The record includes the authority for these decisions, any limits on alternatives, how concerns will be addressed, and what people are told about their options.	☐ Met ☐ Open ☐ N/A	
Asking questions or making a complaint The available channels, responsible office, and response time are clear and accessible.	☐ Met ☐ Open ☐ N/A	
Protection during review The record describes how a complaint or correction request is reviewed, what temporary protection is available, and who is authorized to correct the issue.	☐ Met ☐ Open ☐ N/A	
Complaint and correction outcomes The record identifies who communicates and records the outcome, remedy, and any unresolved action. It links to the completed record for an actual case or, before operation begins, the tested procedure.	☐ Met ☐ Open ☐ N/A	
Participation informed product selection The record includes who participated and what support enabled them to do so. It explains how their questions and recommendations influenced the choice before the product or its setup was fixed.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Participation during local evaluation The record includes the input received while the service was being assessed and how the team responded or changed the proposal.	☐ Met ☐ Open ☐ N/A	
Recommendations still unresolved The record lists unresolved recommendations from affected people or staff. It names the reviewer responsible, explains why each remains unresolved, and states the next action.	☐ Met ☐ Open ☐ N/A	
If involvement was missed If affected people or staff were not involved before product selection or local assurance, the record explains why and how their absence affects the decision. The record states what involvement is required before the next decision.	☐ Met ☐ Open ☐ N/A	
Affected groups and barriers to access The record identifies locally relevant groups and barriers involving language, disability, literacy, digital access, or other circumstances. It provides the baseline, sample, and limitations of the assessment.	☐ Met ☐ Open ☐ N/A	
Equity and accessibility differences assessed The assessment compares each group's result with the reference result and explains any material difference. The record includes the required response or condition for stopping use, the reassessment date, and the accessible alternative.	☐ Met ☐ Open ☐ N/A	
Indigenous governance requirements The named Indigenous governance or Nation-led process determines which requirements apply and who has authority.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Required Indigenous-led process completed The record includes the required protocol, engagement, agreement, or approval.	☐ Met ☐ Open ☐ N/A	
Procedure for patient-brought AI advice Evidence confirms that the procedure names who assesses the concern, checks relevant claims against clinical evidence and the person's circumstances, records the source and care plan, and escalates suspected harm. Consumer AI output must not be treated as verified clinical evidence.	☐ Met ☐ Open ☐ N/A	
Conversations about external AI tested Testing covers respectful discussion of the tool used, information entered, advice received, and action taken. The discussion explains uncertainty, corrects unsafe claims, establishes agreed next steps, and checks understanding. Test scenarios include medication changes, delayed urgent care, and anxiety. Patients need not provide passwords, account access, or their entire chat history.	☐ Met ☐ Open ☐ N/A	
Older adults and diverse needs included Design and testing include older adults and people with different language, disability, health-literacy, and digital-access needs.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Individual support arrangements The assessment establishes how staff will identify individual support needs and provide accessible formats, assistance with use, and a non-digital option. It covers caregiver or proxy involvement in line with the person's wishes and applicable consent and capacity rules. The assessment also considers privacy when devices are shared.	☐ Met ☐ Open ☐ N/A	
Hospital responsibility explained Information distinguishes tools people choose independently from services the hospital recommends, purchases, connects, or deploys. It explains differences in data handling, access to human help, and accountability. Any proposal for hospital adoption, recommendation, integration, or systematic reliance on a tool requires governance review.	☐ Met ☐ Open ☐ N/A	
Care available without a consumer AI account Having a consumer AI account is never a condition of care.	☐ Met ☐ Open ☐ N/A	
Patient-facing safety arrangements tested The record includes tests with intended users showing whether they understand the service, can obtain urgent help, and can reach a person when needed. Tests must also assess whether users can use alternatives when the AI is unavailable or unsafe. The monitoring and readiness record includes the test evidence, limitations, and responsible owner.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Patient-facing risks monitored Monitoring covers misleading reassurance, delayed care, access barriers, and work passed to clinical teams. The record includes evidence, limitations, the responsible owner, review frequency, and response thresholds. These findings are linked to the monitoring and readiness record.	☐ Met ☐ Open ☐ N/A	

The affected decision is Not ready if a required authority, choice process, accessible alternative, or Indigenous-led process is missing. The same applies if a material difference between groups has no acceptable response. The review conclusion appears below, with the reviewer's signed confirmation in **Table L50**.

Supporting file or link:

Review conclusion: ☐ Ready ☐ Ready with required actions ☐ Not ready

Table L50*. Patient and equity review confirmation.*

Reviewer name and role	Signature / electronic confirmation	Date

7.3 Local evidence and evaluation

The evidence reviewer works with clinical, frontline, technical, and operational colleagues to determine whether the findings support the requested use and scope. The review uses the depth assigned in **Section 4**. Results are judged against the approved acceptance criteria, taking account of the workflow and technical setup tested and the people and conditions represented. The finding explains how missing evidence, uncertainty, or failed criteria restrict the uses and settings the hospital can approve.

Table L51 records the local-assurance route selected in **Section 4**. Every route must address the relevant baseline: the configuration and workflow, acceptance criteria, stopping and fallback arrangements, monitoring, changes, incidents, and retirement. Existing evidence may be reused only when it is current and a documented comparison shows that it applies to this use. Streamlined review requires confirmed eligibility and a link to the qualified pathway's central program record and current status; each application does not need a separate program audit. Initial Tier 3 and Tier 4 approvals require full review. A later focused review may address a specific change or question if the reviewer confirms that the remaining baseline and evidence are unchanged and current. This does not make the use eligible for streamlined approval. Full review also requires **Table L52**.

7.3.1 Checks by review route

***Table L51**. Checks for the selected local-assurance route.*

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Use and workflow assessed The record includes the task, users, setting, data, output, resulting action, workflow, and assessment date. This record defines the baseline against which the proposed use will be assessed.	☐ Met ☐ Open	
Technical setup assessed The record includes the provider, model, release, configuration (settings and the other details of the setup), integration, enabled or disabled features, and assessment date. This record defines the technical setup assessed for live use, which becomes the technical production baseline.	☐ Met ☐ Open	
Eligibility for streamlined review Evidence confirms that the use passed the Tier 0 screening and is classified as Tier 1 or Tier 2. It must have no material effect on care, clinically significant access, rights, employment, resource allocation, patient communication, an official record, or a safety-critical workflow.	☐ Met ☐ Open ☐ N/A	
Match to the approved configuration and use Evidence demonstrates that the proposed use matches the approved baseline and has no unresolved specialist concern or blocking issue.	☐ Met ☐ Open ☐ N/A	
Role competency Current evidence shows that the people who request, review, use, support, and monitor the service are competent for their roles.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Decisions eligible for streamlined review The record links to the qualified program's eligibility rules for Tier 1–2 decisions and the groups of uses it audits. Evidence demonstrates that the population and decision in this request are eligible for the pathway.	☐ Met ☐ Open ☐ N/A	
Streamlined review: audit sample The record links to the approved sampling method and the rationale for the chosen sample in the central program record.	☐ Met ☐ Open ☐ N/A	
Streamlined error and acceptance rules The record links to the material-error definition and acceptance criterion approved before audit results were reviewed.	☐ Met ☐ Open ☐ N/A	
Streamlined review: audit responsibility The linked program record names the program's reviewer, audit frequency, approving authority, and approval date.	☐ Met ☐ Open ☐ N/A	
Evidence authorizing the streamlined pathway The record links to the comparison of the proposed pathway with completed full reviews, the findings for each eligible tier and group of uses, and the decision authorizing use of the pathway.	☐ Met ☐ Open ☐ N/A	
Latest streamlined audit The record links to the latest audit period, findings, material errors, and pathway status. Evidence confirms that the audit was current when this decision was made.	☐ Met ☐ Open ☐ N/A	
Streamlined pathway status Evidence confirms that the pathway is active for this decision. If it is suspended, the record links to the finding that led to suspension, the authority, and the date.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Evidence for reopening the pathway Where the pathway has been reactivated, the record links to verified corrective actions, competency checks, and the qualifying audit. The record identifies who authorized reopening and when.	☐ Met ☐ Open ☐ N/A	
Acceptance and stop criteria approved in advance Evidence confirms that each applicable criterion due for this decision was approved before evaluation began. The record states whether it is critical and what response it requires.	☐ Met ☐ Open	
Evaluation coverage justified The record includes the sampling method and counts for relevant populations, subgroups, difficult cases, users, sites, shifts, and workload.	☐ Met ☐ Open	
Release data and processing verified The data and sequence of processing steps intended for release have been verified.	☐ Met ☐ Open	
Evaluation gaps and uncertainty assessed The assessment compares coverage and results with the requirements set before testing. The record explains deviations, how readily errors can be detected, the time needed to correct them, and the remaining uncertainty. It states how these findings limit permitted use or require further evidence.	☐ Met ☐ Open	
Safe alternative workflow defined The record identifies the manual or previous process, when it must be activated, and who may authorize it. It states the maximum volume and duration it can support, available staffing, and the capacity of services receiving the work.	☐ Met ☐ Open	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Safe alternative tested against targets The assessment compares exercise results with downtime and recovery targets. The comparison covers communication, reconciliation of work and records, and the response when demand exceeds capacity.	☐ Met ☐ Open	
Limits and review triggers The record explains the remaining uncertainty. It identifies changes to the model, provider, data, prompt, threshold, interface, integration, user, population, setting, action, degree of independence, or workflow that require a new review.	☐ Met ☐ Open	

7.3.2 Additional checks for full review

***Table L52**. Additional evidence checks for full review.*

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
External evidence assessed The record links to research and supplier evidence relevant to the product, release, intended purpose, and outcome. It identifies local evidence needed where suitable external evidence is absent. A gap cannot count as a favourable finding.	☐ Met ☐ Open	
Local applicability assessed The assessment compares the external evidence with the proposed population, case mix, setting, workflow, users, devices, input quality, timing, and the sequence of data processing and system steps used in live operation. The record explains how important differences affect this decision.	☐ Met ☐ Open	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Comparison and reference outcome The record describes the comparison with the current workflow and relevant human performance. It defines the reference standard or outcome used to judge results, how disagreements are resolved, and what uncertainty remains.	☐ Met ☐ Open	
Sample size and coverage justified The evaluation justifies sample size and event counts for relevant subgroups, rare serious cases, difficult cases, missing inputs, and exclusions. The record explains how gaps limit approval.	☐ Met ☐ Open	
Testing cases outside the main evaluation Testing covers populations, settings, inputs, and missing data that differ from the main evaluation set.	☐ Met ☐ Open ☐ N/A	
Response to unusual cases The report states whether the service recognized an unusual condition and issued a warning or safely deferred the task. The assessment establishes whether a dangerous output could go undetected.	☐ Met ☐ Open ☐ N/A	
People using the AI in practice The assessment covers reliance on outputs, overrides, error severity, decision time, workload, and subsequent actions using realistic roles and working conditions.	☐ Met ☐ Open	
Capacity to act on outputs Test results confirm that the service receiving the work can handle the expected volume within the required time.	☐ Met ☐ Open	
Required human checks feasible and performed Observation confirms that required checks can be and are carried out.	☐ Met ☐ Open	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Missed checks can be detected and addressed Evidence confirms that missed checks and other departures from required practice can be detected and addressed.	☐ Met ☐ Open	

The reviewer must approve the criteria in **Table L53** before evaluation begins. Each criterion has a recorded applicability finding, a critical or non-critical designation, a threshold, and the decision before which it must be met. The result is linked to its units, denominator or sample, reporting period, subgroup rules, and uncertainty so it can be compared with the threshold. Any N/A reason must be specific to the service and accepted by the reviewer; for example, supplier restoration time may not apply to an internally operated service. A failed safety requirement cannot be reclassified as non-critical after the result is known.

***Table L53**. Acceptance criteria and assessment results.*

Criterion, applicability, and N/A reason	Threshold; critical or non-critical; decision before which it must be met	Observed result and relevant subgroup or difficult-case result	Pass, Fail, Insufficient, or N/A; uncertainty	Evidence; approving reviewer and approval date
Minimum acceptable benefit			☐ Pass ☐ Fail ☐ Insufficient ☐ N/A	
Minimum task-performance result			☐ Pass ☐ Fail ☐ Insufficient ☐ N/A	
Maximum acceptable serious errors			☐ Pass ☐ Fail ☐ Insufficient ☐ N/A	
Minimum acceptable result for required subgroups			☐ Pass ☐ Fail ☐ Insufficient ☐ N/A	

Criterion, applicability, and N/A reason	Threshold; critical or non-critical; decision before which it must be met	Observed result and relevant subgroup or difficult-case result	Pass, Fail, Insufficient, or N/A; uncertainty	Evidence; approving reviewer and approval date
Usability acceptance criterion			☐ Pass ☐ Fail ☐ Insufficient ☐ N/A	
Maximum acceptable burden of human review			☐ Pass ☐ Fail ☐ Insufficient ☐ N/A	
Maximum added workload			☐ Pass ☐ Fail ☐ Insufficient ☐ N/A	
Maximum downtime			☐ Pass ☐ Fail ☐ Insufficient ☐ N/A	
Maximum acceptable data loss			☐ Pass ☐ Fail ☐ Insufficient ☐ N/A	
Recovery-time target			☐ Pass ☐ Fail ☐ Insufficient ☐ N/A	
Supplier acknowledgement or response time			☐ Pass ☐ Fail ☐ Insufficient ☐ N/A	

Criterion, applicability, and N/A reason	Threshold; critical or non-critical; decision before which it must be met	Observed result and relevant subgroup or difficult-case result	Pass, Fail, Insufficient, or N/A; uncertainty	Evidence; approving reviewer and approval date
Supplier restoration time			☐ Pass ☐ Fail ☐ Insufficient ☐ N/A	
Required coverage of safety-critical monitoring			☐ Pass ☐ Fail ☐ Insufficient ☐ N/A	
Maximum delay in monitoring data or reports			☐ Pass ☐ Fail ☐ Insufficient ☐ N/A	

Every applicable critical criterion due for this decision must Pass, including required minimum subgroup results and the safe-alternative test. A critical Fail or Insufficient result blocks the proposed scope. N/A requires an accepted reason and cannot set aside a relevant safety requirement. For a non-critical Fail or Insufficient result, the finding explains the effect on expected benefit and remaining risk, names the action and owner, and records why the decision-maker is holding, narrowing, or permitting the use. Missing evidence is not a Pass. Criteria due only at a later decision remain pending until then. The conclusion appears below, with the reviewer's signed confirmation in **Table L54**.

Supporting file or link:

Review conclusion: ☐ Ready ☐ Ready with required actions ☐ Not ready

***Table L54**. Evidence and local-assurance review confirmation.*

Reviewer name and role	Signature / electronic confirmation	Date

7.4 Checks by AI capability

The enabled and requested functions in **Section 1** determine which checks in **Table L55** apply. One product may need several sets because its functions can fail differently and need different safeguards. The relevant specialist decides which checks apply to each component. Each check records how it was assessed, the acceptance rule set

before testing, the result, reviewer, date, and decision requiring the evidence. An unresolved check remains Open, with any necessary restriction recorded until the evidence is accepted.

Predictive or discriminative checks apply to scores, classifications, and rankings. The relevant generative checks apply to drafting, summarizing, answers based on retrieved material, and extraction into structured fields. Ambient capture, transcription, and clinical documentation also need the scribe checks. Systems that use tools or act across other systems need the agentic checks, even when a person confirms each action. Externally managed component checks apply to external hosting, model services supplied by other providers, and dynamic routing. An unknown capability remains Open until resolved. Uncertainty does not justify N/A.

***Table L55**. Checks for predictive, generative, scribe, agentic, and external components.*

Feature	**Status and applicability**	**Required checks for each applicable function**	**Reviewer and date; evidence ID and result** **If Open: interim restriction, action owner, due date, retest, and escalation**
Predictive output and decision rule	☐ N/A ☐ Met ☐ Open	The target outcome or reference result used to judge correctness is defined. The record specifies the score, class, threshold, or ranking rule and the expected user response, with an explanation of the consequences of false positives and false negatives.	
Predictive performance	☐ N/A ☐ Met ☐ Open	Measures suit the task and include discrimination, calibration, and clinically important errors where applicable. Discrimination concerns how well the system distinguishes cases with different outcomes or classes. Calibration concerns agreement between predicted probabilities and observed outcomes. Assessment covers sites, time periods, and subgroups using the approved decision rule. Tests assess whether performance depends on suspected incidental cues (shortcut learning) and examine performance on data from relevant sites or periods excluded from model development.	
Predictive input failures	☐ N/A ☐ Met ☐ Open	Test results show the response to missing, late, outdated, or out-of-range inputs, including how data are delivered in live operation.	
Predictive workflow	☐ N/A ☐ Met ☐ Open	Alert or work volume, user response, overrides, and fallback have been tested.	

Feature	Status and applicability	Required checks for each applicable function	Reviewer and date; evidence ID and result If Open: interim restriction, action owner, due date, retest, and escalation
Changes affecting predictive performance	☐ N/A ☐ Met ☐ Open	The monitoring plan identifies signals for changes in data, outcomes, performance, workflow, and suppliers, and the findings that require recalibration, retraining, rollback, or renewed approval.	
Generative system setup and limits	☐ N/A ☐ Met ☐ Open	The record specifies permitted and prohibited uses, permitted data, model and prompt or template versions, limits on sources and retrieval, settings, logging, and retention.	
Support for generated statements	☐ N/A ☐ Met ☐ Open	Tests assess statements that are unsupported or contradicted by the evidence, using an acceptance rule set in advance.	
Important omissions	☐ N/A ☐ Met ☐ Open	Tests assess clinically or operationally important omissions, using an acceptance rule set in advance.	
Sources for generated answers	☐ N/A ☐ Met ☐ Open	Where the use requires source attribution, tests assess whether references or links are accurate and accessible and whether the system responds safely when a source cannot be verified. Otherwise, the record explains why this check does not apply.	
Refusal and safe deferral	☐ N/A ☐ Met ☐ Open	Tests establish whether the system refuses a request or defers the task when required and provides safe next steps for prohibited, unsupported, or uncertain cases. The findings explain the limits of relying on refusal as a safeguard.	
Variation between generated outputs	☐ N/A ☐ Met ☐ Open	Repeatability is tested across approved settings against defined limits for acceptable variation.	

Feature	Status and applicability	Required checks for each applicable function	Reviewer and date; evidence ID and result If Open: interim restriction, action owner, due date, retest, and escalation
Resistance to unauthorized instructions in inputs or sources	☐ N/A ☐ Met ☐ Open	Tests assess attempts to make the system follow unauthorized instructions through approved inputs and retrieval sources. They include instructions entered directly by a user and instructions hidden in material the system reads (direct and indirect prompt injection).	
Preventing disclosure of protected data	☐ N/A ☐ Met ☐ Open	Tests establish whether the system reveals sensitive information, data held within a tenant's account or environment, prompts, or retrieved material that it should not disclose.	
Required output format	☐ N/A ☐ Met ☐ Open	Where structured output is required, tests cover its fields, required structure (schema), and format, and whether the receiving system interprets it correctly. Otherwise, the record explains why N/A applies.	
Performance in each approved language	☐ N/A ☐ Met ☐ Open	Each approved language is tested against its pre-set acceptance rule.	
Generative accessibility	☐ N/A ☐ Met ☐ Open	The approved output is tested with the required accessibility methods and assistive technology.	
Work needed to verify generated output	☐ N/A ☐ Met ☐ Open	The time and effort required for the assigned human verification step have been measured.	
Patient-facing use and further distribution	☐ N/A ☐ Met ☐ Open	The record describes controls on patient-facing use and on output passed to other people or systems, explaining how they prevent use beyond the approved purpose and limits.	
Generative change control	☐ N/A ☐ Met ☐ Open	The record explains how changes to the model, prompt, template, or reference source are controlled and which changes require repeat testing and renewed approval.	

Feature	Status and applicability	Required checks for each applicable function	Reviewer and date; evidence ID and result If Open: interim restriction, action owner, due date, retest, and escalation
Scribe functions and whether they are enabled	☐ N/A ☐ Met ☐ Open	Every scribe function has an enabled, disabled, or unavailable status in the inventory for the requested release. The record identifies any required human review or authorization for actions, including who may authorize them, and links to supporting evidence.	
Scribe capture and patient link	☐ N/A ☐ Met ☐ Open	The record describes the capture method and whose speech or information is included. Evidence shows how each capture is linked to the correct patient or encounter.	
Scribe data handling	☐ N/A ☐ Met ☐ Open	The record identifies the chart information supplied or retrieved, retention, hosting, other organizations processing data for the provider (subprocessors), and any use beyond the original purpose.	
Scribe version and setup	☐ N/A ☐ Met ☐ Open	The record links to the approved model, release, prompt, template, safeguards, and integration. It states whether the output is a draft or final record and the authorization required for any action, including who may authorize it. It also specifies the specialty, language, device, setting, exclusions, and fallback.	
Scribe capture start and stop	☐ N/A ☐ Met ☐ Open	Visible controls for starting, pausing, resuming, and stopping capture are tested in the authorized environment and verified to work as intended in the workflow being released.	
Scribe patient information	☐ N/A ☐ Met ☐ Open	Tests establish whether the approved information reaches patients and whether they understand it.	
Scribe patient choice	☐ N/A ☐ Met ☐ Open	Tests establish how a patient refuses or opts out and how care proceeds without the scribe.	

Feature	Status and applicability	Required checks for each applicable function	Reviewer and date; evidence ID and result If Open: interim restriction, action owner, due date, retest, and escalation
People who should not be recorded	☐ N/A ☐ Met ☐ Open	Tests establish how people who should not be recorded are excluded. If capture occurs, tests assess how those people are notified and how the recording or resulting records are corrected.	
Identifying speakers correctly	☐ N/A ☐ Met ☐ Open	Tests establish whether the scribe attributes speech to the correct person in the approved languages, devices, and encounter settings.	
Correct patient and encounter	☐ N/A ☐ Met ☐ Open	Tests confirm that every capture and draft is linked to the correct patient and encounter.	
Preventing use of unsigned drafts	☐ N/A ☐ Met ☐ Open	For a scribe that requires clinician review, evidence demonstrates that an unsigned draft cannot become final or trigger a downstream action. Automatic finalization or action requires prior separate, explicit authorization. The review must address the potentially more serious consequences of automatic finalization or action.	
Scribe clinician review and attestation	☐ N/A ☐ Met ☐ Open	For a scribe that requires clinician review, evidence demonstrates that the clinician reviews and signs off on the note before it becomes final or is relied on elsewhere. Governance approval covers only the specified scribe use; other enabled functions need their own authorization.	
Scribe correction and complaint route	☐ N/A ☐ Met ☐ Open	Tests establish how patients or staff question or correct output and how affected records are corrected after a material error. A non-scribe route is included where required.	

Feature	Status and applicability	Required checks for each applicable function	Reviewer and date; evidence ID and result If Open: interim restriction, action owner, due date, retest, and escalation
Assessment of clinically important note errors	☐ N/A ☐ Met ☐ Open	Assessment of clinically important note errors uses an independently selected sample that accounts for error severity and an acceptance rule set before evaluation begins.	
Scribe change testing	☐ N/A ☐ Met ☐ Open	Material changes have been tested against the reference encounters.	
Scribe rollback or safe shutdown testing	☐ N/A ☐ Met ☐ Open	Before release, testing confirms that the service can return to an approved working version or be safely disabled with a functioning alternative. The record includes any technical limitation and the accepted continuity arrangement.	
Agent's goal and authority	☐ N/A ☐ Met ☐ Open	The record defines the approved goal, what starts the agent's work, its degree of independence, permitted tools and actions, prohibited actions, and the points at which responsibility must pass to a person.	
Agent's operating limits	☐ N/A ☐ Met ☐ Open	The record links to the permitted actions, systems, and populations, including the maximum unattended operating time and controls for interrupting and reversing actions.	
Agent characteristics used in tiering	☐ N/A ☐ Met ☐ Open	The record describes how independently the agent acts, how it can affect its environment, how complex its goals are, and the range of tasks it performs. These characteristics inform consequence-based tiering; they are not measures of clinical effectiveness or a validated risk score.	
Agent identity and limited permissions	☐ N/A ☐ Met ☐ Open	Verification confirms that the agent operates under a named service account or identity and its credentials grant only the permissions it needs.	

Feature	Status and applicability	Required checks for each applicable function	Reviewer and date; evidence ID and result If Open: interim restriction, action owner, due date, retest, and escalation
Agent action confirmation and intervention	☐ N/A ☐ Met ☐ Open	Confirmation is required in every case specified by the person or body with approval authority. For actions authorized without prior human confirmation, verification covers the permitted limits and how a person can intervene.	
Limits on an agent's activity	☐ N/A ☐ Met ☐ Open	Defined limits cover the number of steps, time, retry attempts, and volume of activity.	
Duplicate and partial-action prevention	☐ N/A ☐ Met ☐ Open	Tests establish whether the system prevents and detects duplicate or partly completed actions. Tests also assess whether the hospital can safely correct and reconcile such actions and records or recover from their effects.	
Simulation of agent behaviour	☐ N/A ☐ Met ☐ Open	Simulation tests cover the approved goals, tools, limits, and handoffs, including the response to unauthorized instructions introduced through inputs or sources (prompt injection) and malicious output from a tool. Tests establish whether a blocked task leads to safe stopping or authorized help, rather than attempts to bypass the limits.	
Record of the agent's actions	☐ N/A ☐ Met ☐ Open	Evidence shows that records allow the hospital to reconstruct the agent's goals, plans, requests to tools, actions, approvals, overrides, errors, and recovery.	
Separation of an agent's stored memory	☐ N/A ☐ Met ☐ Open	For an agent that retains memory, tests confirm that stored information remains within the approved boundaries for patients, encounters, users, and environments. A design that retains no memory between interactions must be documented and verified as stateless.	

Feature	Status and applicability	Required checks for each applicable function	Reviewer and date; evidence ID and result If Open: interim restriction, action owner, due date, retest, and escalation
Retention and deletion of agent memory	☐ N/A ☐ Met ☐ Open	Where memory is retained, tests cover retention periods, expiry, deletion, and permitted retrieval. Where no memory is retained between interactions (a stateless design), linked evidence supports N/A.	
Attempts to act outside approved limits	☐ N/A ☐ Met ☐ Open	Tests demonstrate detection of attempts to act outside the approved goal, permitted tools, or authorized actions. They assess reward hacking, including attempts to alter evaluation criteria or falsely record a task as complete. Independent records in the receiving system verify completion.	
Detecting unusual agent activity	☐ N/A ☐ Met ☐ Open	Evidence shows that unusual patterns in actions, retries, volume, timing, or errors are detected.	
Changes in the agent's goal or plan	☐ N/A ☐ Met ☐ Open	Evidence shows that departures from the approved goal in the agent's active plan or behaviour are detected.	
Consistency with patient goals	☐ N/A ☐ Met ☐ Open	For actions affecting patients, tests assess how the system detects and handles conflicts with recorded goals or preferences. For other uses, the record explains whether this check applies.	
Monitoring the tools and services the agent uses	☐ N/A ☐ Met ☐ Open	Evidence shows that monitoring detects when tools, models, data sources, or integrations become unavailable, perform worse, or change.	
Emergency stop for the agent	☐ N/A ☐ Met ☐ Open	Emergency suspension, containment time, fallback, and preservation of evidence have been tested for the affected function.	

Feature	Status and applicability	Required checks for each applicable function	Reviewer and date; evidence ID and result If Open: interim restriction, action owner, due date, retest, and escalation
Agent restart requirements	☐ N/A ☐ Met ☐ Open	The record identifies who may authorize restart and who may complete production confirmation, the separate check that the actual system is ready to resume live use. It states the evidence required for each decision.	
Externally managed or dynamically routed model components	☐ N/A ☐ Met ☐ Open	The record identifies externally managed components, their model or service providers, and organizations that process data for them (subprocessors). It includes identifiers showing which release is used or how requests are routed, approved limits on changes, and evidence needed to detect or investigate a substituted component.	
Change control for externally managed components	☐ N/A ☐ Met ☐ Open	The record specifies required advance notice and time for local testing of changes, reference cases, release stages, and monitoring after a change. It includes triggers for renewed approval and rights to defer, reject, disable, roll back, or end the arrangement.	

Ambient capture, transcription, or clinical documentation requires the scribe inventory in **Table L56**. Every function has an enabled, disabled, or unavailable status for the requested release. Each enabled function has a recorded human-review requirement or authority permitting its actions, linked to supporting evidence. A function is not authorized simply because it is included in the product. Any enabled action outside the approved scribe use and limits needs a separate review.

***Table L56**. Scribe functions and action authority.*

Scribe function	Enabled, disabled, or unavailable in this release	Required human review or authorization permitting actions	Evidence or version/setup record
Transcript	☐ Enabled ☐ Disabled ☐ Not available		

Scribe function	Enabled, disabled, or unavailable in this release	Required human review or authorization permitting actions	Evidence or version/setup record
Note sections	☐ Enabled ☐ Disabled ☐ Not available		
Full note draft	☐ Enabled ☐ Disabled ☐ Not available		
Structured fields	☐ Enabled ☐ Disabled ☐ Not available		
Coding suggestions	☐ Enabled ☐ Disabled ☐ Not available		
Patient messages	☐ Enabled ☐ Disabled ☐ Not available		
Referrals	☐ Enabled ☐ Disabled ☐ Not available		
Orders	☐ Enabled ☐ Disabled ☐ Not available		
Tasks	☐ Enabled ☐ Disabled ☐ Not available		
Automatic finalization of the record	☐ Enabled ☐ Disabled ☐ Not available		
Other function (describe)	☐ Enabled ☐ Disabled ☐ Not available		

N/A applies only when the reviewer accepts why the check does not apply to this component or use. A Met finding needs the assessment method, acceptance rule set in advance, result, reviewer, and date supporting the requested decision. An Open finding has a Pending, Failed, Unknown, or Expired status and a recorded affected decision,

temporary restriction, required action, owner, deadline, closure or retesting method, and escalation. An Open check due now prevents a Ready conclusion. A later check blocks earlier work only when that work depends on it. The conclusion appears below, with the reviewer's signed confirmation in **Table L57**.

Supporting file or link:

Review conclusion: ☐ Ready ☐ Ready with required actions ☐ Not ready

Table L57*. AI-function review confirmation.*

Reviewer name and role	Signature / electronic confirmation	Date

8 Readiness and ongoing oversight

8.1 Readiness for live use

The operational owner, end-user leads, and internal technical leads use **Table L58** to assess readiness for the requested release or scope change. The assessment covers staff preparation, workflow, and support services, with clinical, education, usability, and workforce reviewers involved as needed. It examines what staff can do under realistic conditions, their competence for assigned roles, the skills they must retain without AI, and how work will be distributed. Evidence must demonstrate that support, downtime arrangements, governance capacity, and recurring funding can sustain the proposed scale. Each outstanding action must be assigned before the decision that depends on it.

***Table L58**. Readiness checks for the requested operating scope.*

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Limits of operation The operating limits specify the sites, users, shifts, volume, dates, and maximum extent of use covered by the initial release, expansion, material change, or restart.	☐ Met ☐ Open ☐ N/A	
Reasons for selecting this work or service The selection rationale explains the measured need and why the selected work is representative of the intended use. It includes local capacity to support the service and report problems.	☐ Met ☐ Open ☐ N/A	
Operating and stopping criteria are set The operating plan specifies the criteria for continued use, mandatory stop thresholds, and first review date.	☐ Met ☐ Open ☐ N/A	
Access and suspension authorities are assigned The operating plan identifies who is authorized to grant access and who is responsible for emergency suspension.	☐ Met ☐ Open ☐ N/A	
Participants in workflow testing Workflow testing includes representative people who create, review, or act on output, as well as those who take over responsibility for it. Learners and teams receiving the work participate when they are affected.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Workflow scenarios Testing assesses routine work and work under time or workload pressure against defined pass criteria. Scenarios cover disagreement with output; incorrect, late, or missing output; downtime; handover; and escalation. Where patients may bring external AI advice, testing assesses respectful evaluation and correction of that advice.	☐ Met ☐ Open ☐ N/A	
Workflow test results The test record describes the environment and use tested, results, and unresolved actions, with responsible owners and retest dates. The workflow that will actually be released is verified before operational access is granted.	☐ Met ☐ Open ☐ N/A	
Role-specific competency assessed before access Each role's tasks and decisions are linked to learning objectives and a scenario assessment completed before access. The record includes the pass criteria, assessor, result, and date. Assessment covers approved data use, evidence and uncertainty, verification, bias and accessibility, escalation, fallback, and patient communication where relevant. Rotating learners, temporary staff, and new users are included.	☐ Met ☐ Open ☐ N/A	
Education delivery is arranged The education plan specifies the responsible person, delivery method, and protected time, including arrangements for rotating learners, temporary staff, and new users.	☐ Met ☐ Open ☐ N/A	
Access matches competency A linked, dated list of eligible users shows each person's training, assessment, and access status. The record explains any mismatch between the competency required and the access granted.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Required competency reassessments are completed Reassessment is completed after a material change, incident, audit failure, role change, or a prolonged lapse in relevant practice.	☐ Met ☐ Open ☐ N/A	
Failed assessments lead to restricted access or authority Access or authority is restricted after a failed assessment.	☐ Met ☐ Open ☐ N/A	
Skills needed without AI The competency record identifies the skills that must remain available without AI support. It records how they were assessed at baseline, the results, and the responsible owner.	☐ Met ☐ Open ☐ N/A	
Maintaining skills without AI For each required skill, the competency record specifies the minimum acceptable result, reassessment method and frequency, sample assessed, practice needed to maintain the skill, and response if performance falls short.	☐ Met ☐ Open ☐ N/A	
Changes to who does the work The workforce assessment describes work moving between roles, sites, or shifts and records the baseline workload and people consulted.	☐ Met ☐ Open ☐ N/A	
Acceptable workload The workload plan states how workload will be measured, the maximum acceptable burden, and who must respond when that limit is exceeded.	☐ Met ☐ Open ☐ N/A	
Patient and user information is ready Approved information is available for patients and users.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Patient and user support is ready Accessible support, response times, and correction or complaint routes are ready. The responsible team is identified and confirms its readiness.	☐ Met ☐ Open ☐ N/A	
A non-AI alternative is ready where feasible or required The team responsible for the alternative confirms its readiness. The record states any limits and how concerns will be addressed. A missing mandatory alternative prevents release.	☐ Met ☐ Open ☐ N/A	
Downtime and support capacity The continuity plan describes the alternative workflow, when it is activated, and who decides to activate it. It specifies the maximum volume the workflow can handle, how long it can operate, staffing, support hours, and the communication route.	☐ Met ☐ Open ☐ N/A	
Downtime and recovery test result The test record compares results with recovery and data-loss targets. It includes reconciliation of work and records, defects, responsible owners, deadlines, and evidence from retesting.	☐ Met ☐ Open ☐ N/A	
Required and prohibited user practice The user-practice record describes what users must and must not do, names the responsible owner, and explains how adherence will be observed.	☐ Met ☐ Open ☐ N/A	
Results of auditing user practice The audit compares current adherence to required practice with the baseline and threshold and identifies the audit source, sample, reviewer, and date.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
When required practice is not followed The operating plan describes how the interface, access permissions, and supervision support required practice. It states the response to repeated failures to follow that practice, including access restrictions or escalation where required.	☐ Met ☐ Open ☐ N/A	
Capacity to govern this use The linked current governance-capacity assessment covers the proposed use and its scale. It identifies any new backlog, breached threshold, or required response.	☐ Met ☐ Open ☐ N/A	
Fit with the configuration group The record explains how the requested use fits the approved operating limits and funded capacity. It states whether the number of configuration groups would change or an exception is proposed.	☐ Met ☐ Open ☐ N/A	
Ongoing funding The recorded funding covers the proposed use and scale, including monitoring, support, incidents, continuity, and exit. The record is updated if the owner, funded period, or risk to renewal changes.	☐ Met ☐ Open ☐ N/A	

Optional AI learning resources: AI Learning Hub; Start Here; Living Glossary; International AI in Medicine Education Working Group; Global Framework for an AI Curriculum in Medicine

8.2 Monitoring, changes, incidents, and retirement

Table L59 records a measure for each relevant monitoring area. The measure cell contains the reviewer's accepted reason when an area does not apply. Separate rows are required whenever the source, threshold, owner, or response differs. Each entry links to the complete **Appendix I** monitoring record, including the denominator, reporting window, data completeness, uncertainty, subgroup results, and maximum permitted delay. The record also specifies any separate stop rule needed for a rare but severe event, so acceptable overall results do not conceal a severe event that requires action.

***Table L59**. Monitoring measures, thresholds, and owners.*

Concern or outcome	Measure or accepted N/A reason; unit, population, source, and maximum delay	Baseline and target	Warning, action, and stop threshold	Owner, review frequency, and monitoring-record link
Safety and quality				

Concern or outcome	Measure or accepted N/A reason; unit, population, source, and maximum delay	Baseline and target	Warning, action, and stop threshold	Owner, review frequency, and monitoring-record link
Model, output, and data performance				
Equity, access, and patient or user experience				
User practice, overrides, workload, and skill retention				
Reliability, cybersecurity, data flow, and integration				
Value, adoption, and supplier or service performance				

8.2.1 Known and unexpected failures

The review in **Section 8.2** includes the failure coverage checks in **Table L60** and lifecycle readiness checks in **Table L61**. The same hazard assessments and monitoring records can be linked wherever they are needed, without attaching duplicate copies. Met means accepted evidence shows that the check is satisfied for the requested use. Every Open item has a recorded status and linked action in **Section 10.1**.

***Table L60**. Failure coverage and discovery checks.*

Check and required evidence	Result and record
Coverage of foreseeable failures The hazard record identifies applicable failure patterns and links each specific hazard to its safeguard and monitoring record. It includes accepted reasons for exclusions and known detection gaps.	☐ Met ☐ Open Evidence reference and reviewer/date: Open category and action reference:

Check and required evidence	Result and record
Reliable evidence and outcomes The evidence record identifies source records, reference labels or other comparators, and the time before outcomes become available. It explains how disputed labels, missing data, and insufficient samples are handled and states which evidence is obtained independently of the AI system and the processing steps that supply its inputs.	☐ Met ☐ Open Evidence reference and reviewer/date: Open category and action reference:
Review beyond existing alerts Linked random and targeted sampling plans include relevant cases with no warning, relevant excluded cases, abandoned interactions, and rare or underrepresented groups. The plans record exclusions from the review sample and specify the sample size or selection rule, review interval, reviewer, and findings that require escalation.	☐ Met ☐ Open Evidence reference and reviewer/date: Open category and action reference:
Required safety conditions The safety record identifies applicable checks for the correct patient and encounter, required approval, intended destination, and completion of the intended action. It states where independent evidence is obtained, who reviews a breach, and the required response and deadline.	☐ Met ☐ Open Evidence reference and reviewer/date: Open category and action reference:
Unexpected effects and near misses The reporting plan identifies accessible routes for patients and staff, who assesses incoming reports, response deadlines, feedback to reporters, and escalation. It explains how credible but unfamiliar harm is addressed even when routine measures remain within limits.	☐ Met ☐ Open Evidence reference and reviewer/date: Open category and action reference:
Tests of monitoring and response The linked latest authorized exercise assesses whether data are complete and current and checks calculations, notification, response, and fallback. The record identifies faults, responsible owners, retest dates, and restrictions required when evidence or monitoring is unavailable.	☐ Met ☐ Open Evidence reference and reviewer/date: Open category and action reference:

***Table L61**. Change, incident, monitoring, and retirement readiness checks.*

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Detecting model and data changes The monitoring record identifies the information source used to detect changes to data, model, prompt, or threshold and states how often it is reviewed.	☐ Met ☐ Open ☐ N/A	
Detecting service changes The monitoring record identifies the information sources used to detect changes to interfaces, integrations, suppliers or hosting, workflow, permitted actions, or the system's degree of independence. It states how often those sources are reviewed.	☐ Met ☐ Open ☐ N/A	
Reviewing and testing material changes The change review explains which assumptions supporting approval are affected and which specialists must review them again. It records test methods, acceptance criteria set in advance, and decisions. It identifies the production confirmation needed before release to check the actual system's readiness for the changed use.	☐ Met ☐ Open ☐ N/A	
Changes within approved operating limits The review links the approved operating limits and explains which changes to population, setting, interface, data, permissions, version, or workload remain within them. Supporting evidence shows that the variants are equivalent for the use being assessed.	☐ Met ☐ Open ☐ N/A	
Deciding whether drift requires action The drift assessment specifies the periods being compared, minimum sample, and approach to uncertainty. It explains how a suspected change will be confirmed and which findings require an alert or a stop.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Permission to operate during change review The operating plan names who may authorize continued or restricted operation within the existing approval while a change is reviewed. Reaching a stop threshold requires suspension of use and activation of fallback. Any emergency arrangement for continuing the service requires separate authorization, equivalent protection, and an expiry.	☐ Met ☐ Open ☐ N/A	
Review required for a change The review states whether the change stays within approved operating limits, needs a focused review or full reassessment, or creates a separate use. It records any resulting decision about rollback, supplier escalation, new approval, consolidation, or retirement.	☐ Met ☐ Open ☐ N/A	
Incident response deadlines The record links the incident severity definitions and reporting procedure. Each severity level has stated deadlines for acknowledgement, containing harm, investigation, and notification, with a named on-call owner.	☐ Met ☐ Open ☐ N/A	
Preserving incident evidence The preservation plan lists the records and logs that must be preserved, where they are held, who may access them, and who is responsible for preserving them.	☐ Met ☐ Open	
Results of preserving evidence during a test or incident The record describes the latest test or incident result and any gaps, with the person responsible for corrective action and the retest date.	☐ Met ☐ Open	
Supplier incident duties Linked contract provisions establish supplier notice, evidence, investigation, and correction duties, including response deadlines.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Incident decision authority The incident plan identifies who may restrict or stop the service, authorize required notifications, and decide that the incident can be closed. Each responsibility has a named backup.	☐ Met ☐ Open	
Corrections needed before restart The restart record lists actions that must be completed before restart and the evidence and independent verification needed to confirm their completion. Each other follow-up action has an explicit authorization, owner, and deadline.	☐ Met ☐ Open	
Restart criteria The restart criteria state the required results for safety, performance, equity, monitoring, support, and fallback. Restart also requires an authorized decision and production confirmation that the specified release is ready for the permitted use.	☐ Met ☐ Open	
Incident exercise or event results The record identifies the latest exercise or event, result, unresolved action, responsible owner, and retest date.	☐ Met ☐ Open	
Missed incident deadlines The incident plan states when a missed deadline must be escalated, who does so, whom they contact, and the required response time.	☐ Met ☐ Open	
Review schedule The schedule records the first and next review dates, the interval appropriate to the tier, and who receives the report. It states the last date of data included in each review and the maximum reporting delay, with a named owner and backup.	☐ Met ☐ Open ☐ N/A	
Earlier-review triggers and response time The review schedule specifies the findings or changes requiring an earlier review and the time within which it must begin.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Overdue reviews The review procedure states when a missed review deadline must be escalated, the responsible owner, recipient, and required response time.	☐ Met ☐ Open ☐ N/A	
When retirement must be considered Measurable rules specify when use must be reconsidered because of value, safety, performance, equity, changes, supplier problems, or loss of monitoring. The record names the decision owner and review deadline.	☐ Met ☐ Open ☐ N/A	
Retirement plan in place before operation The linked retirement plan identifies who verifies it.	☐ Met ☐ Open ☐ N/A	
Retirement completion is verified when due At retirement, evidence verifies service removal, continuity arrangements, and required data handling, with the verification date recorded.	☐ Met ☐ Open ☐ N/A	
Loss of monitoring For each critical feed or record, the monitoring plan defines when missing, incomplete, or delayed information requires action. It names who detects it and the maximum response time.	☐ Met ☐ Open ☐ N/A	
Response when monitoring fails Reaching a critical monitoring stop threshold requires suspension of the affected function and activation of fallback. The response record includes notifications, decision authority, recovery evidence, and the review date. Continuing the service temporarily requires a separate, time-limited emergency decision and equivalent protection.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Supplier incident-response commitments For each severity level, the record identifies hospital and supplier contacts and the event that starts the response clock. It specifies response and restoration targets, required evidence and decision authority, available remedies, and the escalation deadline.	☐ Met ☐ Open ☐ N/A	
Testing supplier escalation The test record compares results from the latest exercise or event with the required target and identifies shortfalls, corrective actions, owners, and repeat-test dates.	☐ Met ☐ Open ☐ N/A	

Governance approval defines what may proceed and under which conditions. Production confirmation checks the actual release and every prerequisite for live use. Expansion also needs explicit approval of the increased scope. Production confirmation must be repeated before live use after a material release change, expansion, or restart. Staff uptake alone does not establish benefit or safety. Every action due before the requested decision must be closed, including applicable training, support, monitoring, incident response, stopping, fallback, and recovery requirements. The operational readiness conclusion appears below, with the reviewer's signed confirmation in **Table L62**.

Supporting file or link:

Review conclusion: ☐ Ready ☐ Ready with required actions ☐ Not ready

***Table L62**. Operational readiness review confirmation.*

Operational readiness reviewer name and role	Signature / electronic confirmation	Date

8.2.2 Responses to monitoring

Each threshold has a separate row in **Table L63** and a named responder. A warning threshold starts the specified investigation; an action threshold requires the agreed mitigation or restriction. Reaching a stop threshold requires immediate suspension of the affected function, activation of fallback, and incident escalation. The response record states the evidence needed for restart and production confirmation. Emergency continuation requires separate, time-limited authorization and equivalent protection; it cannot waive mandatory duties.

***Table L63**. Threshold responses and restart requirements.*

Trigger or threshold, source, and type: warning, action, or stop	Warning action, owner, and deadline	Required restriction or suspension; person responsible, authority, and deadline	Notification and fallback	Evidence for restart; authorizer; production confirmation

8.2.3 Service and operating changes

Table L64 records drift or other changes in data, model behaviour, configuration, workflow, or supplier service that could affect the reasons for approval. Each entry compares the current finding with the approved baseline and records the decision. The linked monitoring record shows the calculation, period, data completeness, uncertainty, subgroup results, threshold, and required response. For proposed expansion, **Table L65** records the assessment of operating evidence and service capacity.

***Table L64**. Drift and change register.*

Signal or component	Approved baseline and current result	Threshold or classification; required response	Owner, review date, and evidence
Data or population drift			
Outcome or practice drift			

Signal or component	Approved baseline and current result	Threshold or classification; required response	Owner, review date, and evidence
Model or provider change			
Prompt, retrieval, or threshold change			
Workflow or user-practice drift			
Integration, access, or hosting change			
Supplier, service, or contract change			

8.2.4 Evidence and readiness for expansion

***Table L65**. Evidence and capacity for proposed expansion.*

Question and information to provide	Response or evidence reference
Proposed expansion The expansion proposal describes each proposed difference from the approved baseline. It identifies the configuration group, current and proposed number of groups, and what would remain unchanged, with the comparison supporting any claim of equivalence.	
Change classification The review explains whether the expansion fits approved limits, is a material change, or creates a separate use and states which reviews and approvals are needed.	
Evidence from operation For each decision criterion, the record links the baseline, threshold, observation period, overall and subgroup results, incidents, near misses, findings about required user practice, and uncertainty. It explains differences in the use and findings compared with the first operating scope and names the reviewer accepting the evidence.	
Service capacity The capacity assessment compares forecast demand with tested capacity for training, support, work passed to other services, monitoring, supplier response, and governance review.	

Question and information to provide	Response or evidence reference
Portfolio and funding The linked current capacity and funding assessment identifies backlogs or threshold breaches, the number of configuration groups, any funded exception and its expiry, and any required consolidation or retirement.	
Readiness checks The readiness record states the status of integration, contract, privacy, security, fallback, rollback or safe disabling, and tests of stopping the service. Each unresolved limitation has an owner and due date.	
Readiness conclusion	☐ Ready for separate expansion authorization ☐ Not ready Reviewer, date, unresolved limits, owners, and due dates:

9 Legal and commercial review

9.1 Legal, regulatory, professional, and records review

The legal, regulatory, or records reviewer uses **Table L66** to answer the questions assigned in **Section 4**. The review includes professional, billing, accessibility, research, and notification requirements where relevant. It establishes what applies to this product, purpose, and jurisdiction and which licences, approvals, or other authorizations are needed. Identifying a requirement does not grant permission; the record links to the submissions and decisions currently in effect. The affected activity cannot proceed while it is prohibited, its legal permissibility is unresolved, or required authorization is missing. The conclusion appears below, with the reviewer's signed confirmation in **Table L67**.

***Table L66**. Legal, regulatory, professional, and records checks.*

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Regulatory status of the product The regulatory review identifies the jurisdiction, classification, and applicable requirements. Where a licence or authorization is required, it is linked to the exact product, version, and intended purpose. The review identifies any current or proposed use outside that authorization.	☐ Met ☐ Open ☐ N/A	
Professional responsibility The review explains applicable scope-of-practice, accountability, supervision, and documentation requirements, including the professional's duties when relying on or disagreeing with output.	☐ Met ☐ Open ☐ N/A	
Records and documentation The review identifies requirements for official records, documentation, retention, disclosure, and correction and records the reviewer's conclusion.	☐ Met ☐ Open ☐ N/A	
Billing and coding duties are met The review records applicable billing or coding requirements and the reviewer's conclusion, supported by evidence.	☐ Met ☐ Open ☐ N/A	
Human-rights and accessibility duties are met The review records applicable requirements, its conclusion, and supporting evidence.	☐ Met ☐ Open ☐ N/A	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Employment and labour requirements The review identifies applicable requirements, required consultation, and authorization, with its conclusion and supporting evidence.	☐ Met ☐ Open ☐ N/A	
Research duties are met The review records applicable research authority, requirements, its conclusion, and supporting evidence.	☐ Met ☐ Open ☐ N/A	
Public-sector duties are met The review records applicable public-sector requirements, authority, its conclusion, and supporting evidence.	☐ Met ☐ Open ☐ N/A	
Required external authorities The review identifies each regulator, insurer, professional body, or other authority whose involvement is required and names who is responsible for the process with each.	☐ Met ☐ Open ☐ N/A	
Required submissions and notices The review gives the reference, date, and scope of each required submission or notice and identifies anything still outstanding.	☐ Met ☐ Open ☐ N/A	
Required external decisions in effect The review identifies each required decision, its conditions, effective and expiry dates, and who verified that it is in effect.	☐ Met ☐ Open ☐ N/A	
Unresolved legal questions The review describes each unresolved question, advice sought, temporary restriction, responsible owner, and deadline. An activity cannot be authorized while it is prohibited, its legal permissibility is unresolved, or required permission is missing.	☐ Met ☐ Open ☐ N/A	

Supporting file or link:

Review conclusion: ☐ Ready ☐ Ready with required actions ☐ Not ready

***Table L67**. Legal and records review confirmation.*

Legal, regulatory, or records reviewer: name and role	Signature / electronic confirmation	Date

9.2 Procurement, supplier, and contract review

Procurement, vendor-management, and legal reviewers assess agreements for purchased products, partnerships, and external services on which the AI use depends, including arrangements that combine these routes. Their assessment establishes whether the agreements provide the evidence, operational control, and support required for the intended use and whether the hospital can exercise its contractual rights when the product changes or the relationship ends. The component responsibility table in **Section 5** identifies duties that must be enforceable. Signed terms, approved exceptions, unresolved terms, and requirements that do not apply are recorded as separate findings.

Each contract topic in **Table L68** receives a Signed, Exception, Open, or N/A finding based on the evidence available for the activity under review. Signed requires a link to the executed agreement and clause. Exception requires the authorized decision-maker's written decision, the accepted gap, equivalent protection, and review or expiry date. An exception cannot waive a mandatory duty or remove a core control.

An Open finding has a Pending, Failed, Unknown, or Expired status. For an Open item, a Block finding applies if it prevents proof-of-concept (POC) access, purchase commitment, governance approval, or production confirmation. Later applies only if the item is not yet due and has an owner and deadline recorded in **Section 10**. N/A requires an accepted reason why the requirement does not apply.

***Table L68**. Contract requirements and their decision effects.*

Contract requirement	Status, effect, and evidence reference
Intended and prohibited use The agreement must match the approved users, purpose, data, actions, limits, and support arrangements. It must prevent the provider from extending the use on its own.	☐ Signed ☐ Exception ☐ Open ☐ N/A Effect of an Open item on this activity: ☐ Block ☐ Later Agreement and clause, or exception decision with approver and date:
Product documentation and evidence access The agreement provides access to current product documentation and the evidence needed for local evaluation.	☐ Signed ☐ Exception ☐ Open ☐ N/A Effect of an Open item on this activity: ☐ Block ☐ Later Agreement and clause, or exception decision with approver and date:

Contract requirement	Status, effect, and evidence reference
Evaluation, investigation, and retest support The agreement requires timely test support, investigation assistance, fixes, and repeat testing after material change.	☐ Signed ☐ Exception ☐ Open ☐ N/A Effect of an Open item on this activity: ☐ Block ☐ Later Agreement and clause, or exception decision with approver and date:
Service availability and maintenance The contract states the uptime target, advance notice for planned maintenance and notification deadlines for unplanned maintenance, reporting evidence required, owner, and expiry.	☐ Signed ☐ Exception ☐ Open ☐ N/A Effect of an Open item on this activity: ☐ Block ☐ Later Agreement and clause, or exception decision with approver and date:
Incident response and restoration The contract states acknowledgement, response, and restoration times for each severity, including the recovery-time target and maximum acceptable data loss, expressed as a time interval (recovery-point objective).	☐ Signed ☐ Exception ☐ Open ☐ N/A Effect of an Open item on this activity: ☐ Block ☐ Later Agreement and clause, or exception decision with approver and date:
Continuity, escalation, and remedies The agreement identifies the required continuity-test frequency, escalation deadlines, remedies or service credits, responsible owner, and expiry.	☐ Signed ☐ Exception ☐ Open ☐ N/A Effect of an Open item on this activity: ☐ Block ☐ Later Agreement and clause, or exception decision with approver and date:
Change notice and evidence The agreement specifies how the exact release is identified and what advance notice and evidence are required for a change. It states the time allowed for local testing and when a new approval is needed.	☐ Signed ☐ Exception ☐ Open ☐ N/A Effect of an Open item on this activity: ☐ Block ☐ Later Agreement and clause, or exception decision with approver and date:
Hospital control of change The agreement must give the hospital enforceable rights to defer or reject a change, disable the affected function, roll back where feasible, or end the arrangement safely. If the provider cannot keep a version fixed, an acceptable alternative must be agreed. Unapproved substitutions, features, and data uses must remain outside the permitted service.	☐ Signed ☐ Exception ☐ Open ☐ N/A Effect of an Open item on this activity: ☐ Block ☐ Later Agreement and clause, or exception decision with approver and date:

Contract requirement	Status, effect, and evidence reference
Monitoring and audit access The review confirms that the hospital can obtain local monitoring records, logs, version history, reference evidence, and the audit support needed to reconstruct events.	☐ Signed ☐ Exception ☐ Open ☐ N/A Effect of an Open item on this activity: ☐ Block ☐ Later Agreement and clause, or exception decision with approver and date:
Incident investigation and correction The review confirms that the agreement provides timely incident notice, investigation support, corrective and preventive action, verification rights, escalation, and remedies.	☐ Signed ☐ Exception ☐ Open ☐ N/A Effect of an Open item on this activity: ☐ Block ☐ Later Agreement and clause, or exception decision with approver and date:
Data rights and restrictions The agreement specifies contractual controls on inputs, outputs, prompts, recordings, logs, and derived data. These cover use for training or product improvement, resale, further transfers, and deletion.	☐ Signed ☐ Exception ☐ Open ☐ N/A Effect of an Open item on this activity: ☐ Block ☐ Later Agreement and clause, or exception decision with approver and date:
Other companies and jurisdiction The agreement must identify the outside companies the provider uses and their locations. It must address material changes, the hospital's rights to object or terminate, data transfers, and responsibility for compliance.	☐ Signed ☐ Exception ☐ Open ☐ N/A Effect of an Open item on this activity: ☐ Block ☐ Later Agreement and clause, or exception decision with approver and date:
Liability and insurance The agreement records responsibility for harm, breach, failure, inaccurate output, and non-compliance, including all applicable insurance and indemnity requirements.	☐ Signed ☐ Exception ☐ Open ☐ N/A Effect of an Open item on this activity: ☐ Block ☐ Later Agreement and clause, or exception decision with approver and date:
Continuity and exit The agreement specifies the format, timing, and cost of exports and the help required during transition. It covers replacement or fallback, decommissioning, confirmation of deletion, and rights to restrict or terminate the service. Supporting evidence shows that these arrangements will work.	☐ Signed ☐ Exception ☐ Open ☐ N/A Effect of an Open item on this activity: ☐ Block ☐ Later Agreement and clause, or exception decision with approver and date:

Contract requirement	Status, effect, and evidence reference
Supplier contacts for operation and escalation The agreement identifies primary and backup contacts and their roles, including responsibility for keeping contact information current.	☐ Signed ☐ Exception ☐ Open ☐ N/A Effect of an Open item on this activity: ☐ Block ☐ Later Agreement and clause, or exception decision with approver and date:
Responsibilities for each component The review confirms that the signed agreement includes the component responsibility table and makes the assigned duties for evidence, changes, correction, rollback or safe disabling, funding, and exit enforceable.	☐ Signed ☐ Exception ☐ Open ☐ N/A Effect of an Open item on this activity: ☐ Block ☐ Later Agreement and clause, or exception decision with approver and date:

Before production confirmation, the agreement—or an approved arrangement offering equivalent protection—must make the relevant obligations enforceable. These include limits on use, data rights, security, oversight, monitoring, incident response, change notices, audit access, and safe exit. An exception needs written approval before the activity that depends on it; approval cannot make unlawful access permissible, remove a required authorization, or leave a core control without equivalent protection. The review conclusion appears below, with the reviewer's signed confirmation in **Table L69**.

Supporting file or link:

Review conclusion: ☐ Ready ☐ Ready with required actions ☐ Not ready

***Table L69**. Procurement and contract review confirmation.*

Procurement, vendor, or contract reviewer: name and role	Signature / electronic confirmation	Date

10 Decisions and approvals

10.1 Governance decision and conditions

The AI Steering Committee or authorized delegate records the decision in **Table L70**, drawing on specialist findings and other evidence. The decision must stay within the powers recorded in the Working record. **Table L71** records the permitted activity, consequence tier, acceptance of the risk remaining after safeguards, limits and conditions, monitoring arrangements, and next review. The rationale explains failed or uncertain criteria and any dissent. Required permission from an outside body remains Open until authorization is in effect for the activity concerned.

***Table L70**. Governance decision and permitted activity.*

Decision and scope	Recorded selection and details
Decision being considered The decision record identifies the activity under review and the exact use, scale, and limits the decision would cover.	☐ Initial governance approval ☐ Continued use within the current authorization ☐ Change ☐ Restriction ☐ Pause ☐ Restart ☐ Retirement Scope:
Decision outcome The decision record states what is permitted now and every condition. Governance approval does not permit live release. Initial release, a material release change, expansion, or restart also requires production confirmation that the actual system meets the approval conditions and is ready for live use.	☐ Approve the specified activity ☐ Approve with recorded conditions ☐ Defer pending further evidence or action ☐ Reject Activity authorized, limits, effective date, expiry, and further decisions required:

***Table L71**. Governance rationale, approved baseline, and oversight conditions.*

Final tier and remaining risk The record states the Tier 1–4 classification and credible consequences that determine it. It explains the risk remaining after verified safeguards, the uncertainty, and the decision to accept or reject that risk. Accepting the remaining risk does not lower the consequence-based tier.	

Authority to accept the remaining risk The record names the decision-maker, the basis of their authority, the decision date, and review or expiry date.	
Approved technical baseline The baseline identifies the approved product or model, release, configuration (settings and other details of the setup), data, integrations, and enabled features.	
Approved use and exclusions The approval specifies authorized users, population, setting, workflow, actions, sites, scale, and duration and states which uses remain prohibited.	
Decision about configuration groups The decision record links the approved configuration group, approved operating limits, current number of groups, permitted maximum, and capacity assessment. It includes any new decision about an exception, consolidation, or retirement and its conditions.	
Operating support The record confirms that funding and capacity commitments cover the approved use. It identifies any change in owner, funded period, or renewal date, and any response required to a capacity problem.	

Reasons for the decision The rationale explains why the permitted activity is reasonable, considering expected benefit, important review findings, failed or uncertain criteria, and any disagreement.	
Approved monitoring record The record identifies the monitoring record or dashboard and its thresholds, required responses, owners, and reporting route.	
Review schedule and earlier-review triggers The review schedule specifies the first review and routine frequency, including the findings or changes that would require an earlier review.	
External authority where required The record identifies the external authority, finding or approval, reference, use covered, conditions, effective date, and expiry. The required authority must be in effect before the activity it governs begins. A later decision that is still pending remains an explicit condition until it takes effect.	
Assessment against the hospital's risk tolerance ☐ Within delegated tolerance ☐ Escalation required ☐ Approved exception within policy The assessment explains the result against the applicable criterion and policy and names the authorized approver and conditions. Escalation and exceptions cannot waive a prohibition or mandatory duty.	

Decision record and committee quorum The record names the committee or authorized person and gives the meeting or decision-log reference. Where applicable, it records the quorum: the attendance needed for the committee to make a decision.	
Specialist findings and disagreement The record summarizes specialist conclusions and any dissent and identifies mandatory findings that cannot be overridden by an ordinary committee vote.	
Registry and dossier update The record names who updated the registry and dossier and when. It confirms that they show the decision, lifecycle status, approved baseline, configuration group, approved operating limits, conditions, and next review.	
Portfolio and capacity update Configuration group and capacity records reflect this decision. The update records only changes since the sourcing and component responsibility assessment.	
Funding update The update confirms the initial funding commitment or records changes in funding, owner, funded period, renewal date, and contingency arrangements.	
Linked lifecycle records The applicable monitoring, change, incident, restart, expansion, or retirement record is linked.	

Each required action has a separate row in **Table L72**, identifying which decision must wait and what evidence will demonstrate completion. An action due before the requested governance decision requires deferral, rejection, or a revised scope for review; the blocked scope cannot receive approval. An action due only before release may remain open after governance approval if the approved activity does not depend on it. Live use must still wait for completion. During operation, an overdue action triggers the recorded response. A stop-critical overdue action requires suspension of the affected use and activation of fallback while any further decision is sought.

***Table L72**. Actions required for governance and release decisions.*

#	Action required	Decision that must wait, or later decision before which the action is due	Owner and due date	How completion will be checked	Who may authorize closure; who closed the action and when
1					
2					
3					

10.2 Owner acceptance and governance approval

Governance approval requires signed findings from **Sections 5–9**. Every applicable review must conclude Ready or Ready with required actions for the decision requested now. Later verification may remain pending only if the reviewer records which activity must wait and what evidence will demonstrate completion. An earlier review may be reused only after a documented comparison confirms that its findings remain current, effective, and applicable to the use, configuration, data routes, setting, and conditions.

Approval cannot be signed if a prerequisite review or its scope is unclear, a required finding is missing or Not ready, essential external authorization is absent, or an action due before approval remains open. A review that does not apply requires an accepted reason in **Section 4**. Every permitted later action appears in **Section 10.1** with its owner,

deadline, and affected decision. An ordinary committee vote cannot override a mandatory finding within a specialist's authority. Production confirmation is a separate release decision and must verify every prerequisite for live use.

10.2.1 Owner acceptance

By signing **Table L73**, each owner accepts the approved scope and the safeguards, monitoring, escalation, stopping duties, and required actions assigned to their role.

***Table L73**. Owner acceptance signatures.*

Role	Name	Signature or electronic confirmation	Date
Clinical/end-user owner			
Internal technical owner			
Operational owner			
Executive sponsor, if required			

10.2.2 Governance approval signatures

The approver identified in the Working record signs or electronically confirms the decision in **Section 10.1** using **Table L74**. Entering the committee's name does not replace that person's approval. **Table L75** records any required executive or board decision, or the reason it is not required.

***Table L74**. Governance approval signature.*

Role	Name	Signature or electronic confirmation	Date
AI Steering Committee chair or person authorized under local policy			

***Table L75**. Executive or board decision record.*

Executive leadership or board decision, if required Where local policy requires this approval, the record includes the decision, meeting or reference, named approver, date, scope, and required actions. Otherwise, it explains why approval is not required.	

10.3 Production confirmation

Table L76 must be complete before the first live release and before live use following a material release change, expansion, or restart. Its checks verify the actual release and permitted use, staff competency, safeguards, monitoring, support, and fallback. The record links to the governance approval and any separate expansion or restart authorization. Production confirmation checks readiness within that authorized scope; it cannot enlarge it.

Each row is Met only when all applicable parts of the check pass; any part that does not apply needs a recorded, accepted reason. All applicable checks must be Met before release. The release must match its approval. Authorization, competency, monitoring, stopping, incident response, and safe fallback cannot be waived. An Open prerequisite requires Delay or Return for review.

***Table L76**. Production confirmation checks.*

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Match to the approved release Verification confirms that the installed or enabled product, model, release, prompts, reference content, thresholds, and features match the approved baseline or authorized change.	☐ Met ☐ Open	
Match to approved data and technical arrangements Before access for live use is granted, the release environment's data sources, integrations, service identities, access roles, and hosting arrangements are compared with the approved design.	☐ Met ☐ Open	
Match to the permitted use Verification compares users, sites, population, workflow, permitted actions, volume, and maximum scale with the governing approval.	☐ Met ☐ Open	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Conditions completed before release Each condition has recorded evidence of completion and its acceptance, with the verifier, date, decision authority, and owner's acceptance. Any Open condition required for live use prevents production confirmation.	☐ Met ☐ Open	
Current user competency The competency required for each role is verified before access is granted, with the verifier and date recorded.	☐ Met ☐ Open	
Support team readiness The readiness record identifies the support roles, coverage, and escalation route and links evidence that these arrangements are ready.	☐ Met ☐ Open	
Staff, patient, and public information Verification confirms that approved, accessible information is ready for each relevant audience. An accepted reason is recorded where an audience is not affected.	☐ Met ☐ Open	
Monitoring tested before release Testing under authorized production conditions assesses applicable data feeds, data completeness and delay, alerts, acknowledgement of alerts, and reporting. It demonstrates any approved manual monitoring used for the service.	☐ Met ☐ Open	
Stopping and escalation tested Testing under authorized production conditions demonstrates the required controls for suspension, containing harm, escalation, and safe disabling.	☐ Met ☐ Open	
Testing incident communication Testing under authorized production conditions demonstrates the incident route, required notifications, and applicable staff or patient communications.	☐ Met ☐ Open	

Check	Result	Evidence, result, and date If Open: action, owner, due date, decision that cannot proceed, and escalation
Capacity of the safe alternative Testing demonstrates that the manual or alternative process can handle the approved downtime and volume.	☐ Met ☐ Open	
Recovery and reconciliation Testing covers rollback where available, or the approved arrangement for safely disabling the service and recovering. It includes reconciliation of affected records and the support needed for that work, with defects and retest results recorded.	☐ Met ☐ Open	

Production confirmation must wait when a required decision, finding, signature, authorization, or pre-production action is missing, expired, or inconsistent with the proposed release. The relevant competency, monitoring, incident response, suspension, fallback, and recovery arrangements must be tested and accepted before release. Testing itself needs authorized access and does not permit the system's outputs to guide care or operations.

Table L77 records the decision and scope, **Table L78** the first review and record update, and **Table L79** the required signatures.

Table L77*. Production confirmation decision and scope.*

Production confirmation: readiness for live use	Decision or supporting record
Production confirmation decision The decision states whether it concerns initial release, material change, expansion, or restart. Confirmation covers only the recorded release and limits of use. For Delay or Return, the record lists the blocking actions, responsible owners, and next decision date.	☐ Confirm production ☐ Delay ☐ Return for review
Confirmation record The record identifies the type of confirmation, governing decision reference, and exact release. It specifies the sites, population, users, maximum extent and volume of use, activation date and time, first review, and expiry, with unresolved actions and the next decision date.	

Table L78*. First operating review and record update.*

First operating review The record specifies the date of the first review, its accountable owner, and the recipient of the report.	

Registry and dossier update Verification confirms that the records are current for the approved release and use, status, owners, conditions, next review, monitoring link, and decision reference.	

***Table L79**. Production confirmation signatures.*

Role	Name	Signature or electronic confirmation	Date and time
Clinical/end-user owner			
Internal technical owner			
Operational owner			
Named production confirmation approver			

10.4 Expansion approval

An increase in the approved scope requires review of operating evidence and recorded expansion authorization in **Table L80**. A proposal that changes the evidence, risk, or approved operating limits requires review from the affected specialists and a governance decision defining the permitted change. A materially different use requires a separate intake. Expansion authorization defines the new scope; **Section 10.3** production confirmation is still required before the expanded service starts. **Table L81** records the owners' and approver's confirmation.

***Table L80**. Expansion evidence and authorization decision.*

Expansion review and required evidence	Response or evidence reference
Exact expansion requested The proposal specifies additions or changes to site, specialty, population, language, users, volume, integration, features, actions, or the system's degree of independence. It gives the proposed start date and maximum extent of use requested.	
Expansion readiness and configuration group The review links the readiness and ongoing oversight assessment, differences from the approved baseline, evidence of equivalence, number of configuration groups, and funded capacity. It explains whether the proposal stays within approved operating limits, is a material change, or creates a separate use.	
Evidence applies to the expansion For each proposed difference, the review explains whether and how evidence from current operation supports readiness for the expanded use and its ongoing oversight. It states the remaining uncertainty and restrictions and names the reviewer accepting the evidence.	
Capacity at the proposed scale The assessment confirms that capacity for training, support, work passed to other services, monitoring, supplier response, and governance review can meet forecast demand at the expanded scale.	
Portfolio and funding at scale The record confirms the number of configuration groups and funding committed for the proposed scope. It identifies any change in owner, funded period, renewal date, or approved exception.	
Required specialist findings for the expansion The record links the findings needed for the proposed wider use, including integration, contract, privacy, and security reviews. Each open condition has an owner and deadline, with the decision that must wait identified.	

Expansion review and required evidence	Response or evidence reference
Operational safety controls at the expanded scale The record links test evidence for fallback, recovery, and stopping arrangements under the proposed wider use and identifies checks still needed before production confirmation.	
Expansion decision	☐ Authorize the defined expansion subject to production confirmation ☐ Authorize with later conditions ☐ Hold ☐ Return for review ☐ Restrict or pause The record names the approver and the basis of their authority. It states the exact permitted use, conditions, date, expiry, and the limit beyond which another expansion decision is needed. Live use cannot increase until production confirmation for the expanded scope is complete.
Registry and dossier update The update names who changed the registry and dossier and when. It includes the approved scope and conditions, production confirmation reference, activation status, first review, and linked monitoring or change record.	

***Table L81**. Expansion authorization signatures.*

Role	Name	Signature or electronic confirmation	Date
Clinical/end-user owner			
Internal technical owner			
Operational owner			
Expansion approver			

11 Change, incident, and retirement records

Each material or emergency change, incident or near miss, scheduled review, restriction, pause, restart, or retirement needs a new copy of the **Section 11** record. **Table L82** identifies the event and the decision governing it. **Table L83** records changes, **Table L84** incidents and restart, **Table L85** lessons and corrective-action follow-up, and **Table L86** retirement verification. The record links to the detailed investigation file. Research-to-operations transitions use **Section 5.3**. Any required production confirmation in **Section 10.3** must be complete before the affected service is activated.

***Table L82**. Lifecycle event summary and decision.*

Event and approved use The record identifies the event type, date and time, reporter, tracking number, and status. It identifies the approved release and the use covered by its authorization.	
What happened and how it was detected The account describes the sequence of events, including the first known departure from expected operation and how and when it was noticed.	
People and work affected The account describes the affected period, workflow, sites, data, records, and decisions, including the people or groups involved.	
Effects, severity, and urgency The account describes actual effects and credible potential effects, their severity, how quickly harm could occur, and the urgency of response. It identifies groups that may face a different burden or consequence.	
Immediate response The response record explains what was stopped, restricted, restored, communicated, or preserved and names who authorized each action and when.	

Temporary continuation and its limits The record identifies any separately authorized continuation, why it is justified, its limits, equivalent protection, safe alternative, and expiry. Reaching a stop threshold requires suspension. An emergency arrangement for continuing the service requires a specific decision and cannot waive mandatory duties.	
Reviews reopened and temporary limits The record names each review, reviewer, and deadline and states what is prohibited and what may continue temporarily while review takes place.	
Required notifications and communications The record identifies required notices and recipients, with the owner and deadline for each communication.	
Supporting-record links The incident, change, and investigation records are linked, together with the preserved evidence.	
Event decision The recorded decision is Continue within existing permission, Continue with non-blocking actions, Restrict, Pause, Return to approved version, Repeat testing, Reapprove, or Retire. The record identifies who has authority to decide, the reasons, and the effective date. Closing an incident does not authorize restart.	

Permitted use and expiry The decision specifies the use permitted, when permission takes effect, its limits, and when it expires.	
Follow-up and release requirements The record identifies any change to the tier, measures, or thresholds, with the responsible owner and next decision. Production confirmation verifies readiness for live use before a material release change, expansion, or restart is activated.	

11.1 Routine and emergency changes

***Table L83**. Change assessment, release, and emergency arrangements.*

Required record	Response / evidence
Changes to the technology and its setup The assessment compares the current and proposed model, provider, release, data, prompt or template, threshold, integration, tools, and permissions.	
Changes to the approved use The assessment describes changes to the workflow, users, population, authority to take actions, monitoring, and supplier arrangement.	
Configuration groups and capacity The record identifies the configuration group, approved maximum, current and proposed number of groups, and funded capacity. It includes any exception and its expiry, or required consolidation or retirement.	

Required record	Response / evidence
Parts of the approved configuration that remain unchanged The assessment identifies the baseline elements that must remain fixed through the change.	
Change rationale The rationale explains why the change is needed and how it is expected to affect operations.	
Effect on tier and hazards The assessment explains how the change affects hazards, the tier, and any authority required to accept the risk that remains after safeguards.	
Reopened specialist reviews The record identifies each reopened review, its reviewer, and due date, including any restriction that applies while review proceeds.	
Change-test results The record summarizes results of acceptance tests, tests for unintended effects on previously accepted functions (regression tests), and workflow tests. It gives the evidence identifiers and any unresolved failure.	
Approval to release the change The record links the approved change, permitted release sequence, and production confirmation. It names the authorizer and gives the effective date and conditions that must be met before activation.	

Required record	Response / evidence
Rollback or safe disabling The recovery plan specifies when the response is triggered, the responsible owner, and how service will be recovered. It records the approved version and setup to restore during rollback, or how the affected service will be safely disabled while continuity is maintained. The test result and supporting evidence are linked.	
Monitoring after the change The plan describes additional monitoring measures, responsible owners, and how long the additional monitoring will continue. It gives the first review date and report recipient.	
Emergency procedure The record explains why waiting would have posed greater harm and names who authorized the procedure. It states the permitted use, start time, equivalent protection, expiry, notifications, preserved evidence, and review after the event. It identifies which release and production confirmation checks were completed under the emergency procedure. Required authority cannot be waived.	

11.2 Incidents, corrective actions, and restart

***Table L84**. Incident investigation, correction, and restart.*

Required record	Response / evidence
Affected system and detection The record identifies the affected AI feature and approved baseline. It describes how the issue was detected, the period and people or records affected, its severity, and any gap in monitoring.	
Investigation The record links preserved evidence and the completed investigation. It distinguishes confirmed and contributing causes from unresolved questions and records required notices, investigator, and approval date.	
Containment and correction The response record describes immediate steps taken to contain harm and correct affected records or actions, with the responsible owner and date.	
Preventing a recurrence The corrective plan explains how supplier, policy, workflow, training, or technical changes address the problem and states how their effectiveness will be checked.	
Evidence for restart The review compares evidence with the approved criteria. It includes the status of restored monitoring, support, and fallback, remaining uncertainty, and the decision to accept any remaining risk. It lists any action that still prevents restart.	

Required record	Response / evidence
Restart authorization and confirmation The record names the authorizer and gives the date, permitted use, conditions, and expiry. Production confirmation that the system is ready for live use is linked before the affected function resumes.	
First post-restart review The review plan identifies the owner, review date, evidence period, and measures required for the first review after restart.	

11.3 Incident investigation and closure

***Table L85**. Incident lessons and corrective-action follow-up.*

Required record	Response / evidence
Lessons and record changes The account explains what was learned from the event and links resulting changes to the registry, policy, checklist, contract, training, monitoring, or design.	
Responsibility for corrective action and verification The record identifies the action owner and due date, the verifier, and the method for confirming completion.	
Decision to share learning The decision states whether and where the lesson or incident will be shared and records the authority, intended audience, and date.	

11.4 Retirement verification

***Table L86**. Retirement completion and verification.*

Required record	Response / evidence
Transition and communication The transition record describes the replacement or fallback, service continuity, and required notices, including correction of affected records or actions.	
Access and integration removal The record confirms that accounts, credentials, connection endpoints, tools, and integrations have been removed.	
Data, backups, and logs The record states what was returned, retained, or deleted under the approved plan and links evidence verifying completion.	
Contractual duties and unauthorized copies The record confirms completion of all contractual duties required to end the service and identifies any continuing retention duties. Verification confirms that no unauthorized copies or running instances of the service remain.	
Owner confirmation and retirement record The owner's final confirmation is recorded, and the registry and dossier show the service as Retired.	

Required record	Response / evidence
Retirement lessons, follow-up, and sharing The record summarizes lessons from retirement and the sharing decision and identifies the owner and date for any follow-up.	

Each follow-up action has a separate row in **Table L87**. The Action cell identifies which decision must wait for completion, allowing actions to be assigned and closed independently. **Table L88** records the affected owners' and authorized decision-makers' confirmation, including production confirmation when required.

***Table L87**. Lifecycle follow-up action register.*

#	Action	Owner and due date	How completion will be checked	Who closed the action and when
1				
2				

***Table L88**. Lifecycle decision and confirmation signatures.*

Role	Name	Signature or electronic confirmation	Date and time
Affected owner(s)			
Governance decision-maker			
Production confirmation approver, if needed			