BUSINESS CASE

Checkpoint.

Who decides what, and when?

Checkpoint AI: An AI-supported readmission-risk and follow-up process for Lakeside Health Network, with human decision rights built in

Challenge focus area: Focus Area 2, AI and the Future of Work

Case: Mapping the Human Decision Points, Lakeside General and Riverside Specialist Hospital merger

Prepared by: Pod Nova

Prototype: https://lenavon.github.io/10alytics-Hackathon---Checkpoint-/

Event: 10Alytics Q3 Students Hackathon 2026

Contents

#	Section	Purpose
1	Executive summary	The recommendation in brief
2	Strategic context and case for change	Why act now, and on what problem
3	Options analysis	Alternatives considered and why one is recommended
4	Recommended solution	The TO-BE process and the human-AI decision model
5	Stakeholder analysis and governance	Who is affected, and who decides what
6	Benefits realisation	What improves, for whom, and how it is measured
7	Cost estimate	Illustrative one-time and ongoing costs
8	Financial summary	Illustrative payback and return
9	Risk assessment	Key risks and mitigations
10	Implementation roadmap	Phased delivery plan
11	Success criteria and KPIs	How we will know it worked
12	Assumptions, constraints and dependencies	
13	Recommendation and next steps	
A to E	Appendices	BRD, use cases, traceability matrix, change-impact detail, glossary

1. Executive summary

Lakeside General and Riverside Specialist have merged into Lakeside Health Network and want to use AI to predict which patients are likely to be readmitted within 30 days of discharge, so nurses can prioritise follow-up. The two hospitals run different systems and discharge processes, patients may hold different identifiers at each site, and no one has yet decided who owns a decision once an AI tool is introduced.

The business problem: without a defined decision model, an AI readmission tool risks being built on unreliable, unmatched data, making recommendations no one is accountable for, and losing the trust of the nurses who are meant to act on it.

Recommended option: Introduce the AI readmission-risk tool with a redesigned workflow and four mandatory human-in-the-loop checkpoints (Option C, Section 3), rather than a manual-only fix (Option B) or full automation (Option D).

Case summary

Dimension	Summary
Problem	Two merged hospitals, inconsistent processes, duplicate patient identities, no decision-rights model for AI
Recommended solution	AI-scored, human-approved discharge follow-up workflow with 4 checkpoints (HC1 to HC4) and a 3-level automation policy
Illustrative one-time cost	CAD 1.09 million (Section 7)
Illustrative annual net benefit	CAD 1.18 million per year after ongoing costs (Section 8)
Illustrative payback	Under 12 months (Section 8)
Primary risk	Automation bias and low nurse trust; mitigated by mandatory review, explainability and override rights
Timeframe	28 weeks to supervised go-live, then phased scale-up (Section 10)

2. Strategic context and case for change

2.1 Business problem

The NOVA Analysts will analyse the existing discharge and follow-up process, design the future AI-enabled process, and define the requirements needed to build and implement it safely. Left unresolved, the issues below compound each other: inconsistent data undermines the model, an undefined decision model erodes accountability, and low trust means the tool sits unused even if it is technically accurate.

Issue	Root cause	Consequence if unaddressed
Different processes and systems across the two hospitals	Independent legacy workflows and record systems	Inconsistent model inputs; unequal care standards
Duplicate patient records and different patient IDs	Site-specific identifiers, no network-wide link	Scores built on partial history; patient safety risk
Unclear ownership of decisions after AI implementation	No decision-rights model exists	Diffused accountability; unsafe automation creep
Uncertainty over which AI decisions can be automated	No risk-based automation criteria	Either over-automation of clinical calls, or no automation benefit realised at all
No defined human-review points	Tool not yet designed into the workflow	Harmful AI outputs reach patients unchecked
Risk that staff do not trust the system	Opaque scores, no accountability, no feedback loop	Tool is ignored; investment produces no benefit

2.2 Current state (AS-IS)

Full AS-IS swimlane is provided in Section 4 alongside the TO-BE design, so the change is visible side by side. In brief: discharge documentation, patient identity and follow-up prioritisation are all handled separately by each hospital today, with no shared risk view and no closed feedback loop from readmissions back into practice (see Section 4.1 for the full pain-point list).

2.3 Strategic alignment

Alignment	How this case supports it
Focus Area 2: AI and the Future of Work	Redesigns three jobs (follow-up nurse, Health Records steward, ward staff) around a defined human-AI decision split, not around replacing staff
SDG 3, Good Health and Well-Being	Better-targeted follow-up for patients most likely to be readmitted
SDG 10, Reduced Inequalities	One network identity and process so care does not depend on which hospital holds the record
SDG 17, Partnerships for the Goals	Structured, governed data sharing between two merged institutions
Lakeside's own objectives (BO-1 to BO-5, Appendix A)	Improve follow-up targeting, unify patient records, assign clear decision ownership, use AI safely, and support staff adoption

3. Options analysis

Four options were considered against four criteria: patient safety and trust, cost to deliver, speed to benefit, and fit with the brief's requirement for defined human oversight.

Option	Description	Pros	Cons	Recommended?
A. Do nothing	Continue with manual, judgement-based follow-up prioritisation at both sites	No cost or disruption	High-risk patients continue to be missed; the merger's process gaps persist; no answer to the brief's challenge	No
B. Harmonise process only, no AI	Standardise discharge and identity processes across both sites; keep prioritisation manual	Lower cost; fixes some data problems; low technical risk	Does not use AI to target follow-up; misses the case's core ask and its benefit case	No
C.Checkpoint AI: AI-supported, human-in-the-loop (recommended)	AI scores and tiers readmission risk; humans confirm identity, review low-confidence cases, approve or override every tier, and govern the model monthly	Targets follow-up effort where it matters; keeps accountability with clinicians; builds trust through transparency and override rights; matches the brief exactly	Requires investment in data quality, integration, training and governance	Yes
D. Full automation of prioritisation and scheduling	AI assigns tiers and books follow-up with no routine human review	Lowest ongoing staff effort	No safeguard against data errors or model drift; brief explicitly requires human-review points; highest patient-safety and trust risk	No

3.1 Why Option C

- It is the only option that both uses AI to target follow-up and satisfies the brief's requirement for defined human-review points.
- It keeps clinical accountability with named people (Section 5), which Option D removes and Option B never provides for an AI-driven service.
- Its costs are proportionate to the benefit case (Sections 7 and 8), unlike Option D's undifferentiated automation risk.
- It is deliverable in phases (Section 10), so early learning from a shadow pilot reduces risk before wider rollout.

4. Recommended solution

4.1 AS-IS process (current state, for comparison)

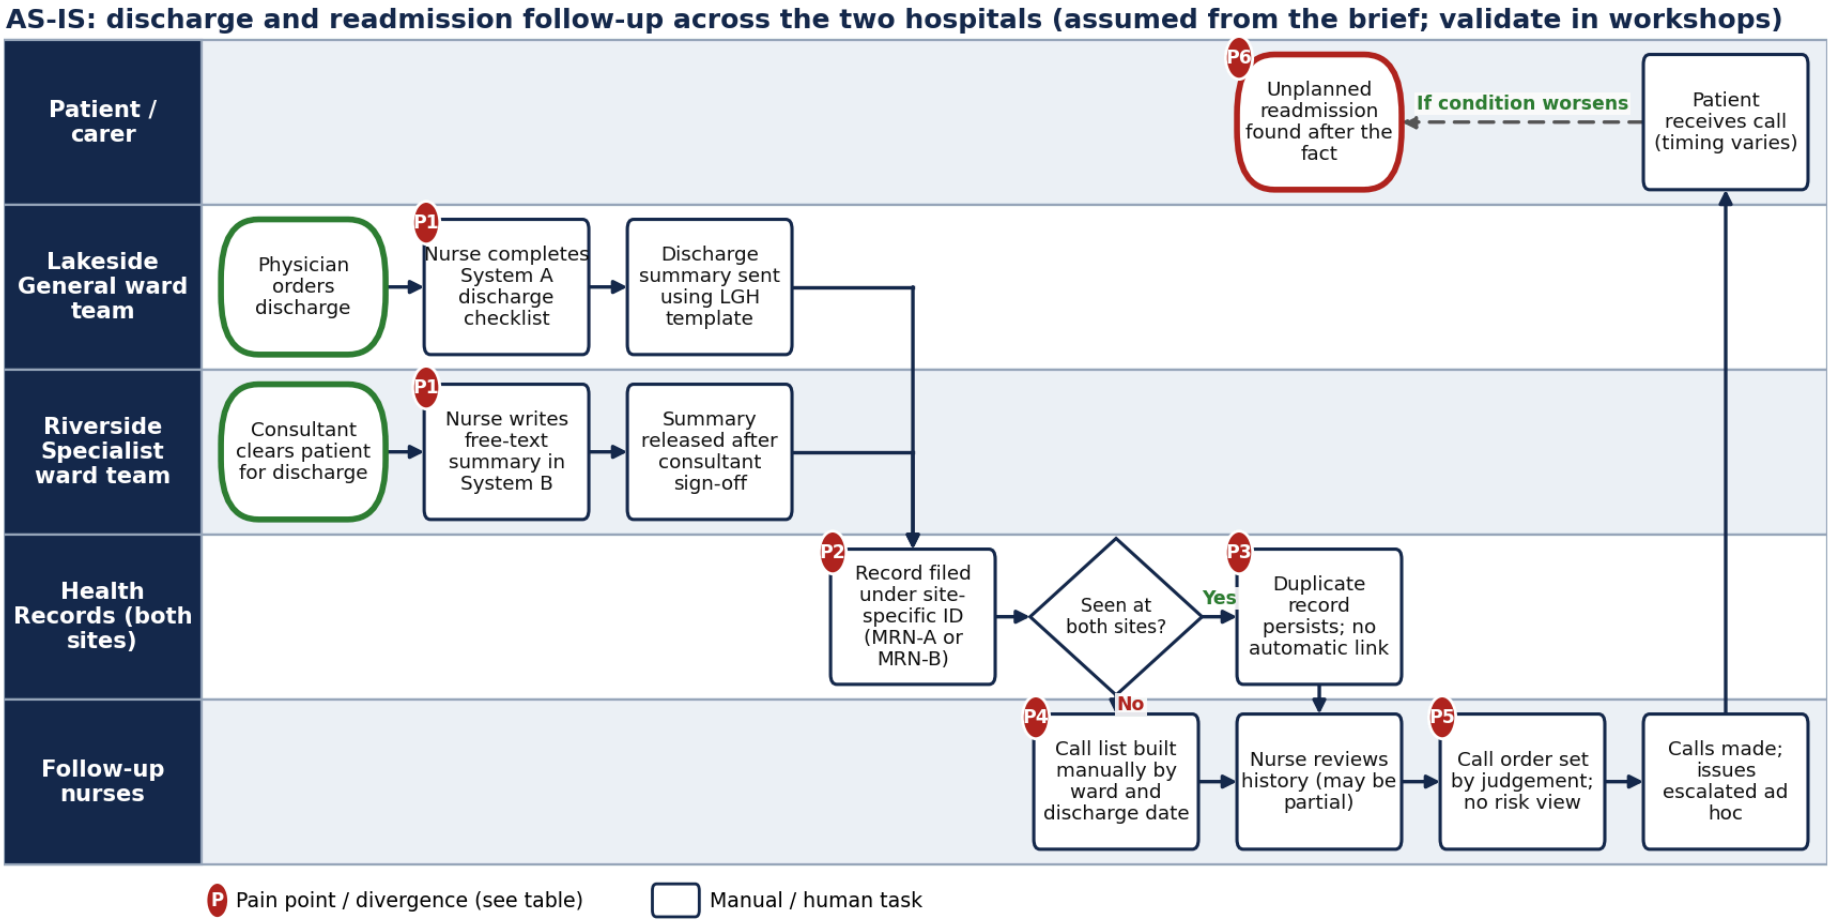


Figure 1. AS-IS swimlane: discharge and readmission follow-up across Lakeside General and Riverside Specialist (assumed from the brief; validate in workshops).

4.2 Pain points driving the recommendation

Ref	Where	What diverges or fails	Consequence for an AI tool
P1	Discharge documentation	System A structured checklist versus System B free-text summary	Model inputs are inconsistent and partly unstructured across sites
P2	Patient identity	Site-specific IDs (MRN-A, MRN-B)	One patient can appear as two people
P3	Duplicate records	No automatic linkage when a patient used both hospitals	Score built on partial history is unreliable, and can understate risk
P4	Follow-up list	Built manually from ward lists by date	Effort spent on assembly, not on patients
P5	Prioritisation	Call order by individual judgement, no shared risk view	Inconsistent care; high-risk patients may wait
P6	Outcome detection	Readmissions found after the fact	No feedback to improve prioritisation

4.3 TO-BE process (recommended design)

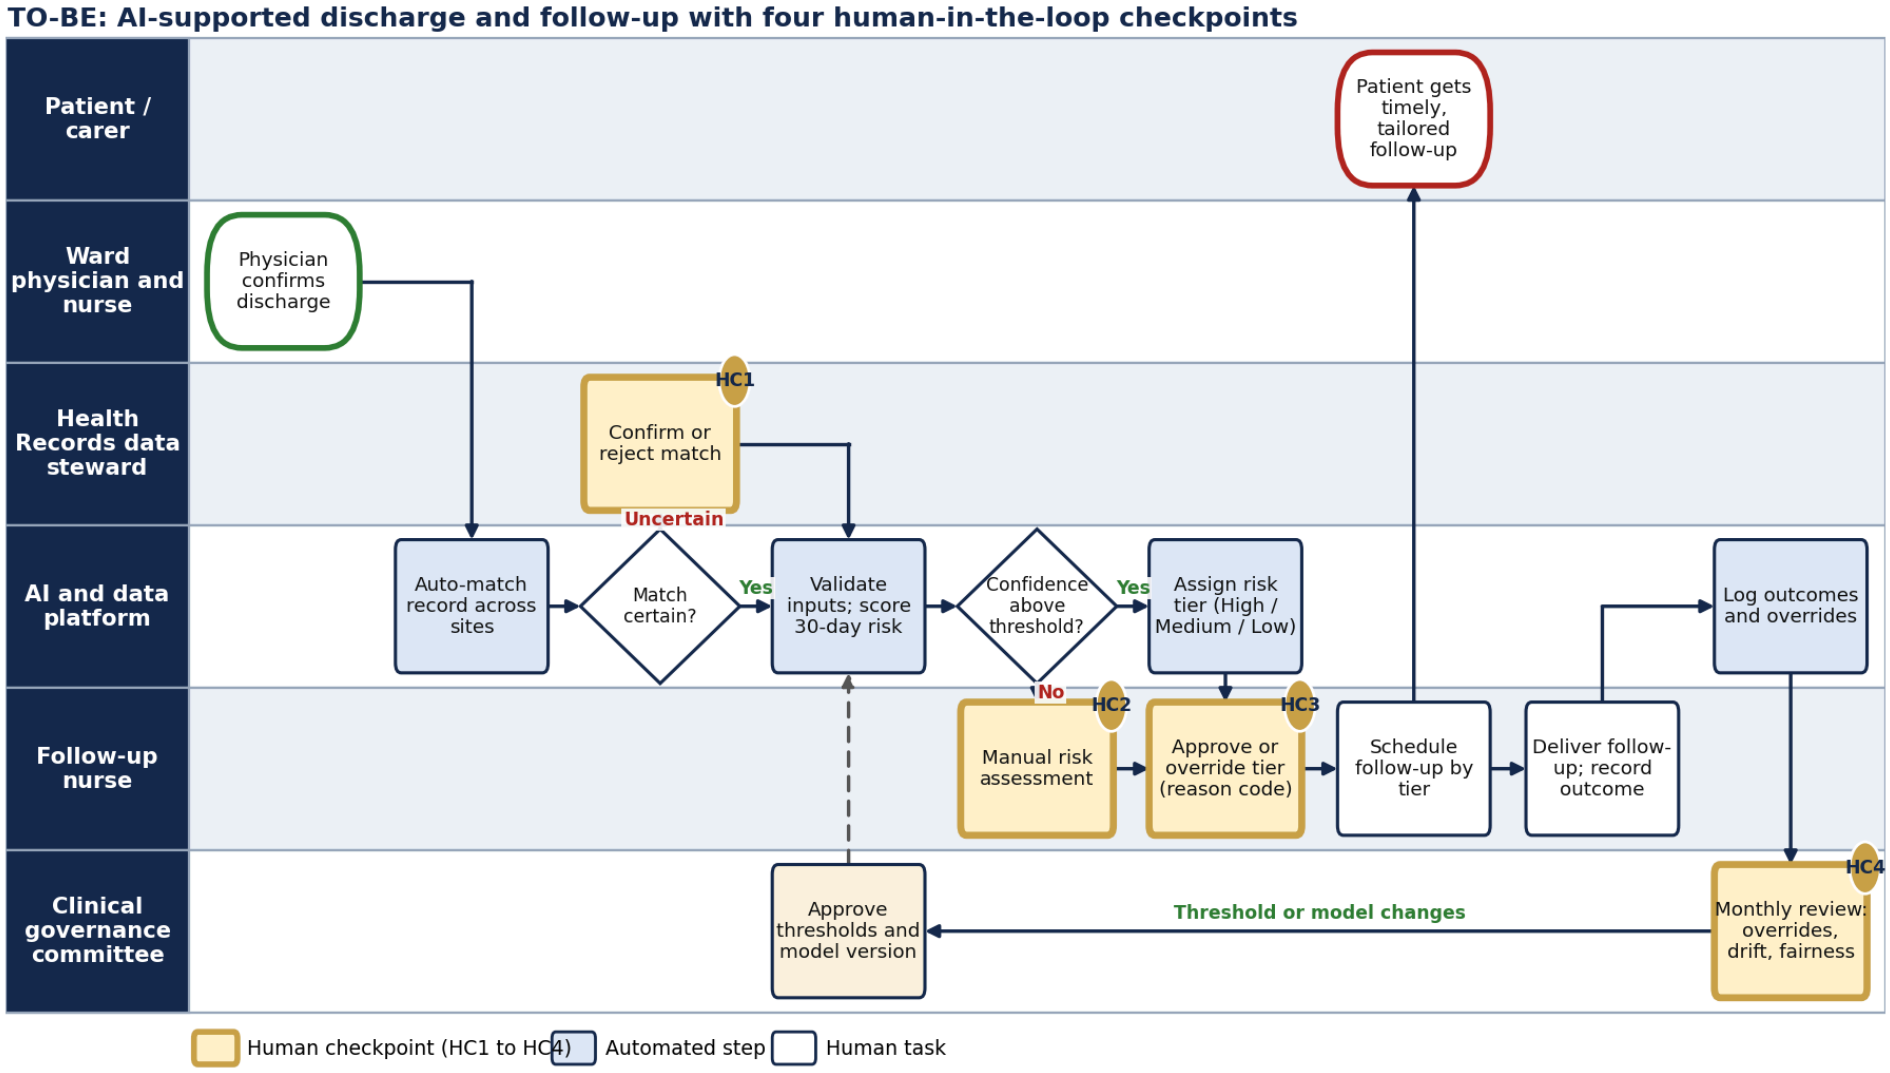


Figure 2. TO-BE swimlane: AI-supported workflow with four human-in-the-loop checkpoints (HC1 to HC4).

4.4 Human-in-the-loop checkpoints

Checkpoint	Trigger	Human role	Allowed actions	Why a human
HC1 Identity match	Match rules return an uncertain match	Health Records data steward	Confirm, reject or request more information; merges are reversible	Wrong merge mixes two patients' histories
HC2 Low-confidence case	Data incomplete or model confidence below threshold	Follow-up nurse	Complete manual risk assessment and set tier	AI should not guess when evidence is thin
HC3 Tier review	Every tier assigned (High approval mandatory; others reviewable)	Follow-up nurse	Approve, or override up or down with a reason code	Clinical accountability stays with a person
HC4 Governance review	Monthly, and on any safety event	Clinical governance committee	Adjust thresholds, retrain, pause the tool	Oversight of drift, fairness and safety

Design principles

AI narrows, humans own. Uncertainty escalates. Every override is a signal for HC4. Downtime never blocks care.

4.5 Human-AI decision framework (governance model)

The brief asks for clear documentation of which decisions AI can make independently and which require human judgement. We use a three-level policy, applied decision by decision.

Level	Meaning	Applies when
A Automate	System acts; result is logged and reversible	Rule-based or statistical, no direct effect on a patient's care, easy to reverse, validated
B AI recommends, human approves	System proposes; a named person must confirm before anything happens	Affects care priorities or patient records, and errors are correctable
C Human decides	A person decides; AI may show information but not choose	Clinical action, lowering care intensity, or any change to how the AI itself works

#	Decision	Level	Accountable role	Rationale
D1	Check required inputs are present before scoring	A	Clinical informatics lead	Deterministic rule; failure routes to HC2, so nothing is hidden
D2	Link records on an exact multi-identifier match	A	Health Records	Deterministic, reversible and logged
D3	Link records on a probable (not exact) match	B	Health Records data steward	A wrong merge mixes two patients' histories (HC1)
D4	Calculate the risk score and its main drivers	A	Clinical informatics lead	Analytical output only; does not itself trigger care
D5	Assign a tier when data is complete and confidence is above threshold	B	Follow-up nurse	AI proposes; nurse approves High and can override any tier (HC3)
D6	Assign a tier when data is incomplete or confidence is low	C	Follow-up nurse	AI must not guess when evidence is thin (HC2)
D7	Set follow-up intensity for High-tier patients	B	Follow-up nurse	Affects workload and patient care; needs a person's agreement
D8	Reduce follow-up below standard care	C	Follow-up nurse	Never automatic; the risk of harm is asymmetric
D9	Clinical action: medication change, physician escalation, readmission	C	Treating physician	Clinical judgement; the tool does not touch these decisions
D10	Change thresholds, retrain or release a model version	C	Clinical governance committee	Changes the behaviour of the whole system (HC4)

5. Stakeholder analysis and governance

5.1 Stakeholder register

Stakeholder	Interest and concern	Influence	Engagement approach
Board of Directors (sponsor)	Safer transitions of care; nurse adoption; accountability	High	Sponsor; approves scope and requirements; monthly steering
Clinical governance committee	Patient safety, fairness, model oversight	High	Owns thresholds and model release; receives monthly dashboard
Clinical informatics / AI lead	Model validity, integration, technical feasibility	High	Co-design; owns model performance monitoring
Follow-up nurses	Workload, trust in scores, autonomy to override	Medium	Primary users; co-design worklist; super-user network
Ward nurses and physicians	Consistent discharge steps; no extra documentation burden	Medium	Workshops on the harmonised discharge standard
Discharge coordinators	Clear handoff to follow-up; timing of summaries	Medium	Process walkthroughs; pilot feedback
Health Records team	Duplicate clean-up workload; role change; job security	Medium	Early involvement; role redesign; training (Appendix D)
IT and integration	Two systems to connect; downtime; security	Medium	Technical design authority; integration plan
Privacy officer	Lawful use of health data; minimum necessary access	High	Privacy impact assessment before pilot
Patients and carers	Timely follow-up; transparency; fair treatment	Low (formal)	Plain-language notice; feedback channel; patient advisory input
Health authority / regulator	Compliance; reporting of readmission indicators	High	Inform; align reporting definitions
Finance and operations	Readmission cost; capacity for follow-up calls	Medium	Benefits tracking; staffing model

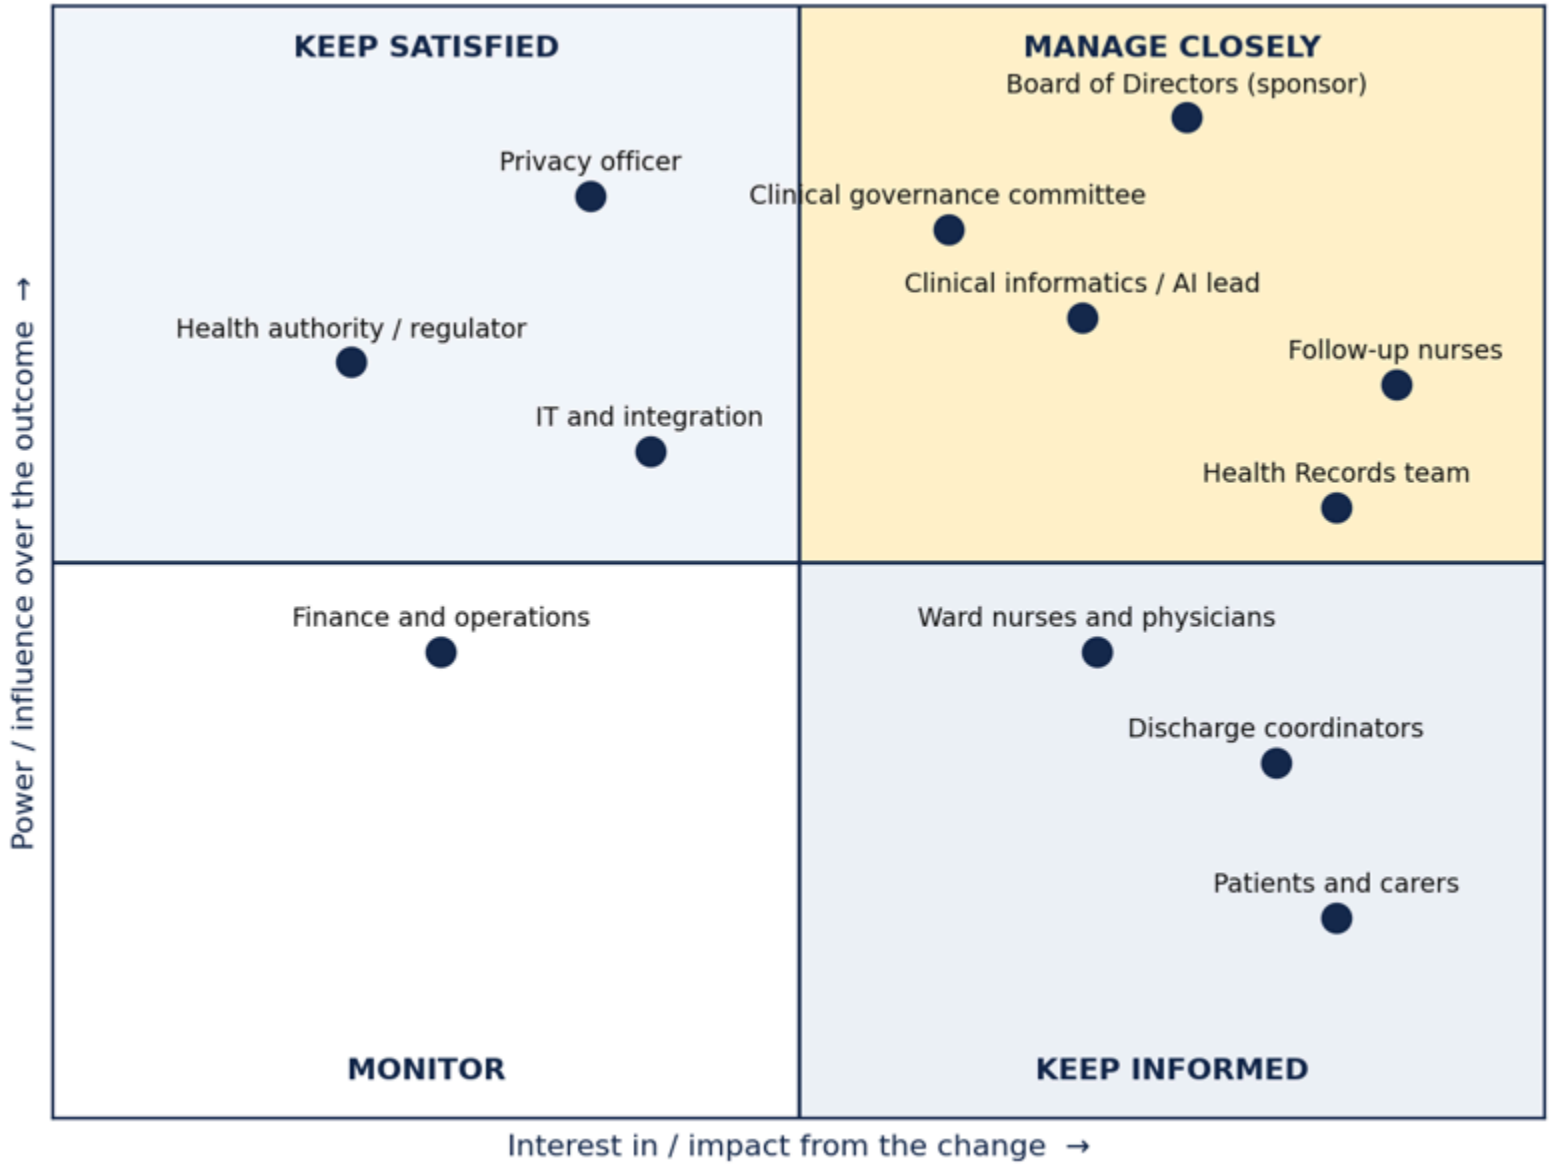


Figure 3. Stakeholder power and interest (positions are analyst judgement, to be validated with the sponsor).

5.2 RACI matrix

R = Responsible (does the work), A = Accountable (one owner who signs off), C = Consulted, I = Informed. Every activity has exactly one accountable role.

Activity	Board of DIrectors	Clinical informatics / AI lead	Ward nurse / physician	Follow-up nurse	Health Records	IT / integration	Privacy officer	Governance committee	Business analyst
Harmonise the discharge standard across both sites	A	C	R	C	C	I	I	C	R
Resolve duplicate records and unify patient ID	I	C	I	I	A/R	R	C	I	C
Set risk thresholds and tier definitions	C	R	C	C	I	I	C	A	C
Generate risk score at discharge (system function)	I	A	I	I	I	R	I	I	I
Review, approve or override the risk tier	I	I	C	A/R	I	I	I	I	I
Schedule and deliver follow-up	I	I	I	A/R	C	I	I	I	I
Monitor accuracy, drift and fairness	I	R	I	C	C	R	C	A	C
Privacy and security impact assessment	I	C	I	I	C	C	A/R	I	C
Training and transition support	A	C	C	C	C	I	I	I	R
Requirements sign-off	A	R	C	C	C	C	C	C	R

6. Benefits realisation

All figures in this and the following two sections are illustrative order-of-magnitude estimates, built to show the shape of the business case. They are clearly marked and would be replaced with Lakeside's own data (readmission volumes, cost per readmission, staff time studies) before an investment decision.

6.1 Tangible benefits (illustrative)

Benefit	How it arises	Illustrative annual value	Linked objecti ve
Avoided readmission costs	Earlier, better-targeted follow-up reduces 30-day readmissions among High-tier patients	CAD 1,200,000	BO-1
Nurse time saved building call lists	Prioritised worklist replaces manual list assembly (est. 30 min/day per nurse across the follow-up team)	CAD 140,000	BO-1, BO-5
Health Records rework avoided	Fewer duplicate-record corrections once identity is resolved up front	CAD 60,000	BO-2
Total illustrative tangible benefit		CAD 1,400,000	

6.2 Intangible benefits

- Clearer accountability for AI-influenced decisions, reducing clinical and reputational risk.
- Higher nurse trust and confidence in the tool, through visible reasoning and override rights.
- A stronger evidence base for future AI initiatives across the merged network.
- A defined career path for Health Records staff moving into data stewardship.

6.3 Benefit ownership

The Board of Directors is accountable for realising the readmission-avoidance benefit; the follow-up nurse lead owns the time-saving benefit; the Health Records lead owns the rework-avoidance benefit. All three report progress to the governance committee alongside the KPIs in Section 11.

7. Cost estimate

Costs are presented as rough order-of-magnitude (ROM) estimates in Canadian dollars, for the purpose of comparing options and sizing the investment. They should be refined with IT, vendor and HR input before funding approval.

7.1 One-time implementation costs (illustrative)

Cost item	Illustrative estimate (CAD)
Patient identity resolution and data clean-up (Health Records, Section 10 Phase 1)	180,000
AI model development, validation and integration	420,000
Workflow and interface build (prioritised worklist, steward queue)	260,000
Change management and training (all affected roles)	140,000
Governance set-up and compliance (privacy impact assessment, security review)	90,000
Total one-time cost (illustrative)	1,090,000

7.2 Ongoing annual costs (illustrative)

Cost item	Illustrative estimate (CAD per year)
Hosting, model monitoring and re-training	95,000
Support and maintenance	120,000
Governance committee time and audit	60,000
Total ongoing cost (illustrative, per year)	275,000

8. Financial summary

This is a simple, illustrative view to support the option decision. It is not a full net-present-value analysis and does not assume a discount rate; Lakeside finance would need to apply one for a funding submission.

Year	Illustrative benefit (CAD)	Illustrative cost (CAD)	Illustrative net (CAD)	Cumulative (CAD)
Year 1	1,400,000	1,090,000 + 275,000 = 1,365,000	35,000	35,000
Year 2	1,400,000	275,000	1,125,000	1,160,000
Year 3	1,400,000	275,000	1,125,000	2,285,000

Illustrative measure	Value
Payback period	Approximately 11 to 12 months from go-live
3-year cumulative net benefit	CAD 2,285,000
Simple 3-year return on the one-time investment	Approximately 210%

Reading this financial case

The readmission-avoidance figure is the most sensitive assumption and the most important to validate first (actual readmission volumes, cost per readmission, and a realistic reduction rate from a shadow pilot).

Even if the readmission benefit is materially lower than illustrated, the efficiency and rework-avoidance benefits alone materially offset the ongoing cost, which supports proceeding to a shadow pilot before a full funding commitment.

9. Risk assessment

Risk	Likelihood	Impact	Mitigation	Owner
Duplicate or wrongly merged records produce a misleading score	Medium	High	HC1, reversible merges, data clean-up (Section 10), match-rate monitoring	Health Records
Over-reliance on the score (automation bias)	Medium	High	HC3 requires review; show confidence and drivers; training on judgement	Board of Directors
Nurses distrust the tool and ignore it	Medium	High	Plain-language drivers, override with reasons, co-design, shadow pilot	Follow-up nurse lead
Model performs unevenly across patient groups	Medium	High	Subgroup monitoring; governance review (HC4)	Governance committee
Privacy breach or unlawful data use	Low	High	Privacy impact assessment; minimum-necessary access; audit	Privacy officer
Model drift after go-live	Medium	Medium	Calibration tracking; version control and rollback	Clinical informatics
Integration delays across two systems	High	Medium	Phase by ward; manual fallback; early technical design	IT
Change fatigue after the merger	High	Medium	Sequenced change; visible sponsor; super-users	Board of Directors
Illustrative benefit estimates overstate the real case	Medium	Medium	Validate readmission and time-saving assumptions in the shadow pilot before wider funding	Business analyst

10. Implementation roadmap

Phase	Timing (illustrative)	Key activities	Exit criterion
0 Prepare	Weeks 1 to 4	Impact validation workshops; baseline duplicate rate; privacy impact assessment	Baseline and role design agreed
1 Clean up and train	Weeks 5 to 12	Reconcile priority duplicates; role-based training; sandbox practice	Duplicate rate at agreed level; training complete
2 Shadow pilot	Weeks 13 to 20	Steward queue live; AI scores produced but not acted on; compare to actual outcomes	Shadow report accepted by governance; benefit assumptions validated
3 Supervised go-live	Weeks 21 to 28	One ward per site; super-user support; full checkpoint model live	Go decision at governance review
4 Scale and embed	Week 29 onward	Extend by ward; move to steady-state monitoring and KPIs	Steady-state queue age and KPIs within agreed limits

11. Success criteria and KPIs

Targets are set only after the Phase 0 baseline. Values shown as illustrative are placeholders, not results.

Measure	Type	Illustrative target	Review cycle	Source
High-tier patients contacted within 2 working days of discharge	Process	≥90% of High-tier patients	Monthly from Phase 2	Worklist data
30-day readmission rate, High-tier patients	Outcome	≥15% relative reduction vs. Phase 0 baseline, within 6 months of go-live	Quarterly	Outcome records
Override rate against AI-proposed tier	Trust	Between 10% and 20% of tiers overridden; any single reason code above 10% of overrides triggers review	Monthly at HC4	Audit log
Duplicate patient record rate	Data quality	From an illustrative 8% baseline to under 2%, within 90 days of Phase 1	Monthly during Phase 1, then quarterly	Health Records
Time to prepare the follow-up list	Efficiency	From ~30 minutes/day per nurse to under 5 minutes/day	At pilot go-live (Phase 3)	Time study
Nurse trust and usability score	Adoption	≥80% of follow-up nurses score the tool 4/5 or higher (post-pilot survey)	End of Phase 2 shadow pilot	Survey
Subgroup performance gap (age, sex, site)	Fairness	Model accuracy varies by no more than 5 percentage points across subgroups	Monthly at HC4	Dashboard
Realised benefit vs. the CAD 1.4M illustrative annual case	Financial	Tracked variance within ±20% of the illustrative figure	Quarterly	Finance

12. Assumptions, constraints and dependencies

Assumptions

- The AS-IS process is assumed from the brief and will be validated in workshops with both sites.
- A validated readmission-risk model is available or can be built; model development itself is technically out of scope of this business case's BA work.
- A data-sharing arrangement between the two sites already exists post-merger.
- Applicable health-privacy legislation is identified by Lakeside's privacy officer; this document does not assume a specific jurisdiction.
- All costs, benefits and timings are illustrative and will be recalibrated with Lakeside data before funding approval.

Constraints

- Delivery must not disrupt live discharge activity at either site.
- Solution must comply with Lakeside's existing security and privacy policies.
- Training capacity is limited during the post-merger transition period.

Dependencies

- Executive sponsorship from the Board of Directors
- IT capacity to integrate the two legacy record systems.
- Availability of Health Records staff for the clean-up phase (Section 10, Phase 1).

13. Recommendation and next steps

We recommend proceeding with Option C: Checkpoint AI An AI-supported readmission-risk and follow-up workflow, delivered through the phased roadmap in Section 10, starting with the Phase 0 preparation and baseline work.

Immediate next steps

- Secure Board of Directors sponsorship and confirm the governance committee's terms of reference.
- Commission the Phase 0 baseline (duplicate rate, current follow-up timing, readmission rate) to replace the illustrative figures in Sections 6 to 8.
- Begin the privacy impact assessment in parallel, so it does not delay the shadow pilot.
- Approve the requirements package in Appendix A as the basis for solution design.

Appendix A. Business Requirements Document

A.1 Purpose and scope

This BRD defines what Lakeside Health Network needs from an AI-supported 30-day readmission-risk process and the changes to work around it. In scope: patient identity across the two sites, risk scoring at discharge, tiering, the follow-up worklist, human review, monitoring and transition. Out of scope: the design of the predictive model itself, changes to clinical treatment pathways, and patient-facing apps.

A.2 Business objectives

ID	Objective
BO-1	Improve follow-up of patients at highest risk of 30-day readmission
BO-2	Establish one trusted view of each patient across both hospitals
BO-3	Give every decision in discharge and follow-up a clear owner and defined human review
BO-4	Use AI safely, fairly and lawfully, with transparency and audit
BO-5	Enable staff to adopt the tool with confidence

Priorities use MoSCoW (Must, Should). Each requirement carries a business objective, traced fully in Appendix C.

A.3 Business requirements

ID	Requirement	Priority	Objective
BR-01	A single harmonised discharge and follow-up process applies at both sites	Must	BO-3
BR-02	Follow-up effort is prioritised toward patients with the highest predicted 30-day readmission risk	Must	BO-1
BR-03	Every AI-influenced decision has one named accountable role	Must	BO-3
BR-04	The solution complies with applicable health-privacy law and Lakeside policy	Must	BO-4
BR-05	A baseline and success measures are agreed before go-live	Should	BO-1

A.4 Stakeholder requirements

ID	Requirement	Priority	Objective
SR-01	Nurses see the main reasons behind each score in plain language	Must	BO-5
SR-02	Nurses can override any tier inside the worklist, with a reason code	Must	BO-3
SR-03	Health Records has a work queue for uncertain matches, with merge and reject tools	Must	BO-2
SR-04	The governance committee receives a monthly report on overrides, performance and fairness	Must	BO-4
SR-05	Patients are told that AI supports follow-up prioritisation and that a person makes the decision	Should	BO-4

A.5 Functional requirements

ID	Requirement	Priority	Objective
FR-01	Link patient records across both sites under a network-wide patient identifier	Must	BO-2
FR-02	Auto-link only on deterministic match rules; route probable matches to the steward queue	Must	BO-2
FR-03	Every merge is reversible and audit-logged	Must	BO-2
FR-04	Check required inputs before scoring; if missing, show 'insufficient data' instead of a score	Must	BO-4
FR-05	Generate the 30-day readmission risk score when discharge is ordered and refresh at discharge	Must	BO-1
FR-06	Display the top contributing factors and a confidence indicator with each score	Must	BO-4
FR-07	Assign High, Medium or Low tier using governance-approved thresholds	Must	BO-1
FR-08	Route low-confidence or incomplete cases to manual nurse assessment	Must	BO-4
FR-09	Present a prioritised follow-up worklist to follow-up nurses	Must	BO-1
FR-10	Require approve or override at HC3; overrides need a reason code	Must	BO-3
FR-11	Never lower follow-up below standard care without a recorded human decision	Must	BO-3
FR-12	Record follow-up outcomes and readmission events to close the feedback loop	Should	BO-1
FR-13	Log every score, tier, view, override and approval with user, time and model version	Must	BO-4
FR-14	Provide a governance dashboard: tier mix, overrides, calibration and subgroup results	Should	BO-4

A.6 Non-functional requirements

ID	Requirement	Priority	Objective
NFR-01	Role-based access, encryption in transit and at rest, and minimum necessary data	Must	BO-4
NFR-02	Logs are tamper-evident and kept per Lakeside retention policy	Must	BO-4
NFR-03	Score available within a proposed 60 seconds of a discharge order (illustrative)	Should	BO-5
NFR-04	Tool downtime never blocks discharge; a manual fallback exists (proposed availability 99.5% in clinical hours, illustrative)	Must	BO-5
NFR-05	Explanations are understandable to a non-technical nurse	Must	BO-5
NFR-06	Performance is monitored across subgroups (for example age, sex, site); disparities beyond an agreed limit trigger review	Must	BO-4
NFR-07	Match rate and data completeness are monitored and reported	Should	BO-2
NFR-08	Approve or override takes no more than three clicks inside the existing worklist (proposed)	Should	BO-5
NFR-09	Model versions are controlled: validated before release, with rollback	Must	BO-4

A.7 Transition requirements

ID	Requirement	Priority	Objective
TR-01	Reconcile and clean legacy duplicate records before go-live	Must	BO-2
TR-02	Role-based training for follow-up nurses, Health Records, ward staff and coordinators	Must	BO-5
TR-03	Run a shadow-mode pilot at both sites: scores produced but not acted on, then compared to outcomes	Must	BO-4
TR-04	Phased rollout: one ward per site first, then wider	Should	BO-5
TR-05	Redefine Health Records roles toward data stewardship	Must	BO-2
TR-06	Support model: super-users, help channel and feedback route	Should	BO-5
TR-07	Documented manual contingency procedures	Must	BO-5
TR-08	Monitor adoption, overrides, safety and outcomes after go-live	Must	BO-5

A.8 Business Requirements Protocol

A.8.1 Unreliable Output Controls

- Restrict AI functions to approved data and validated use cases.
- Do not allow fabricated patient facts to be treated as evidence.
- Display relevant source/context information for material inputs.
- Flag outputs that conflict with source data or validation rules.
- Require human review before defined clinical action.
- Log material AI outputs and review outcomes.

A.8.2 Conflict Resolution

- Identity conflict → stop processing and route to Health Records.
- Clinical-data conflict → flag discrepancy and require clinician review.
- AI recommendation conflicts with professional judgement → human judgement prevails; record override where required.
- AI/system unavailable → use approved fallback workflow.

A.8.3 Compliance Checking

- Validate role-based access before displaying patient information.
- Include privacy/security controls in acceptance testing.
- Maintain an auditable record of material AI-assisted decisions.
- Assign clear ownership for AI monitoring and escalation.

Appendix B. Detailed use cases

UC-01 Prioritise follow-up for a discharged patient

Field	Detail
Primary actor	Follow-up nurse
Supporting actors	Ward physician; AI and data platform; Health Records data steward (only if identity is uncertain)
Trigger	Physician confirms a discharge order
Preconditions	Patient has an admission record at either site Network patient ID is linked, or HC1 has been resolved Nurse has role-based access to the worklist
Main flow	1. Physician confirms the discharge order. 2. Platform checks the network record; the match is certain (FR-01, FR-02). 3. Platform validates inputs and generates the 30-day risk score with top drivers and confidence (FR-04 to FR-06). 4. Platform proposes a tier from the approved thresholds (FR-07). 5. Patient appears in the follow-up worklist, ordered by tier (FR-09). 6. Nurse opens the patient and reads the drivers and confidence. 7. Nurse approves the tier at HC3 (FR-10). 8. Nurse schedules follow-up by tier. 9. Nurse delivers the follow-up and records the outcome (FR-12). 10. Platform logs every step (FR-13).
Alternate flows	A1 Insufficient data (step 3): platform shows 'insufficient data'; case goes to HC2; nurse assesses manually. A2 Low confidence (step 4): routed to HC2, same handling as A1. A3 Nurse disagrees (step 7): nurse overrides the tier and selects a reason code; the override is logged for HC4. A4 Tool unavailable: discharge proceeds on the manual process; missed patients are scored when service returns. A5 Uncertain match (step 2): follow UC-02; if discharge is imminent, the nurse uses manual assessment (HC2).
Exceptions	E1 Nurse attempts to reduce follow-up below standard care without a recorded decision: system blocks the action. E2 Scoring service times out: treated as A4.
Postconditions	Follow-up tier and owner recorded; audit trail complete; outcome captured for the feedback loop
Requirements	BR-02, BR-03, SR-01, SR-02, SR-05, FR-04 to FR-13, NFR-01 to NFR-05, NFR-08, TR-03, TR-07

UC-02 Resolve a possible duplicate patient record

Field	Detail
Primary actor	Health Records data steward
Supporting actors	AI and data platform; follow-up nurse (informed of outcome)
Trigger	Match rules return an uncertain match at discharge or during clean-up

Field	Detail
Preconditions	Records from both sites are readable by the platform Steward has role-based access to the queue
Main flow	1. Platform flags the probable duplicate and places it in the steward queue with the matching evidence (FR-02, SR-03). 2. Steward opens the pair and compares identifiers, dates of birth, contact details, encounters and the match rationale. 3. Steward verifies with the ward or the patient's contacts if needed. 4. Steward confirms the match at HC1. 5. Platform links the records under one network ID and keeps the original site IDs (FR-01, FR-03). 6. Platform re-scores the patient if a follow-up decision is pending. 7. Platform logs the action with steward, time and rationale (FR-13).
Alternate flows	A1 Steward rejects: records stay separate and are marked 'reviewed, not a duplicate'. A2 Not enough evidence: steward requests more information; ward uses manual assessment meanwhile. A3 Error found later: steward unmerges; platform restores both records and re-scores. A4 Time-critical discharge: follow-up nurse proceeds with manual assessment (HC2) while the case is worked.
Exceptions	E1 Clinical data conflicts (for example differing allergy records): escalate to the clinical safety officer before linking.
Postconditions	The patient has one verified network identity, or a documented decision to keep records separate; audit trail complete
Requirements	BR-03, BR-04, SR-03, FR-01 to FR-03, FR-13, NFR-01, NFR-07, TR-01, TR-05

Appendix C. Requirements Traceability Matrix

Each requirement traces upward to a business objective and downward to a process step, a use case and a verification method. Process steps refer to Figure 2 (W1 discharge order; S1 to S4 automated steps; HC1 to HC4 checkpoints; F1 and F2 follow-up).

Req	Objective	Process step	Use case	Verification
BR-01	BO-3	Whole flow	UC-01	Process walkthrough signed off at both sites
BR-02	BO-1	S3, F1	UC-01	Pilot: share of High-tier patients contacted within the agreed window
BR-03	BO-3	HC1 to HC4	UC-01, UC-02	RACI and decision table approved by sponsor
BR-04	BO-4	All	UC-01, UC-02	Privacy impact assessment approved before pilot
BR-05	BO-1	HC4	n/a	Baseline report accepted by governance
SR-01	BO-5	S2, HC3	UC-01	Usability test with follow-up nurses
SR-02	BO-3	HC3	UC-01	Functional test; audit log check
SR-03	BO-2	HC1	UC-02	Functional test with seeded duplicates
SR-04	BO-4	HC4	n/a	Report produced and reviewed during pilot
SR-05	BO-4	F2	UC-01	Patient notice reviewed by privacy officer
FR-01	BO-2	S1	UC-02	Match test on a known duplicate set
FR-02	BO-2	S1, HC1	UC-02	Rule test cases (exact versus probable)
FR-03	BO-2	HC1	UC-02	Unmerge test; log inspection
FR-04	BO-4	S2	UC-01	Test with incomplete records
FR-05	BO-1	S2	UC-01	Timing and refresh test
FR-06	BO-4	S2	UC-01	Display test; nurse comprehension check
FR-07	BO-1	S3	UC-01	Threshold configuration test
FR-08	BO-4	HC2	UC-01	Routing test at boundary values
FR-09	BO-1	F1	UC-01	Worklist ordering test
FR-10	BO-3	HC3	UC-01	Workflow test; step cannot be skipped
FR-11	BO-3	HC3	UC-01	Negative test: no automatic downgrade
FR-12	BO-1	S4	UC-01	Data flow test
FR-13	BO-4	All	UC-01, UC-02	Audit trail review
FR-14	BO-4	HC4	n/a	Dashboard acceptance test
NFR-01	BO-4	All	UC-01, UC-02	Security review and access tests
NFR-02	BO-4	S4	UC-01	Audit test
NFR-03	BO-5	S2	UC-01	Performance test
NFR-04	BO-5	W1	UC-01	Failover drill
NFR-05	BO-5	HC3	UC-01	Usability test
NFR-06	BO-4	HC4	n/a	Subgroup report

Req	Objective	Process step	Use case	Verification
NFR-07	BO-2	S1, S2	UC-02	Data quality report
NFR-08	BO-5	HC3	UC-01	Usability timing test
NFR-09	BO-4	HC4	n/a	Release checklist
TR-01	BO-2	HC1	UC-02	Duplicate rate before versus after
TR-02	BO-5	n/a	n/a	Training completion and competency check
TR-03	BO-4	S2	UC-01	Shadow-mode report
TR-04	BO-5	n/a	n/a	Go or no-go review per phase
TR-05	BO-2	HC1	UC-02	Approved role descriptions
TR-06	BO-5	n/a	n/a	Support log review
TR-07	BO-5	W1	UC-01	Contingency drill

C.1 Coverage check by objective

Objective	Requirements	IDs
BO-1	6	BR-02, BR-05, FR-05, FR-07, FR-09, FR-12
BO-2	7	SR-03, FR-01, FR-02, FR-03, NFR-07, TR-01, TR-05
BO-3	5	BR-01, BR-03, SR-02, FR-10, FR-11
BO-4	13	BR-04, SR-04, SR-05, FR-04, FR-06, FR-08, FR-13, FR-14, NFR-01, NFR-02, NFR-06, NFR-09, TR-03
BO-5	9	SR-01, NFR-03, NFR-04, NFR-05, NFR-08, TR-02, TR-04, TR-06, TR-07

Every objective is covered by at least one Must requirement, and every requirement traces to an objective.

Appendix D. Change impact and transition assessment: Health Records team

The Health Records team is affected more than any other group, because the tool only works if patient identity is right. The work does not disappear; it moves from filing and retrieval to identity stewardship.

D.1 Impact by area

Area	Today	Future	Impact	Mitigation
Core purpose	Register, file and retrieve; correct records on request	Steward the network patient identity; decide uncertain matches (HC1)	High	Redesigned role description, career path (TR-05)
Workload profile	Steady, mostly reactive	One-off clean-up backlog, then queue-based work	High (temporary)	Time-boxed clean-up with temporary capacity (TR-01)
Tools	Two site-specific record systems	One steward queue with side-by-side evidence view	Medium	Sandbox practice and hands-on training (TR-02)
Decision rights	Case-by-case corrections	Written authority to confirm, reject and unmerge, with escalation	Medium	RACI and decision table; escalation path agreed
Performance measures	Volume of records processed	Queue age, match accuracy, unmerge rate	Medium	Co-design measures with the team
Job security concern	Not applicable	Possible fear that automation replaces the team	High (perception)	Sponsor message that the tool creates stewardship work; team involved in design

D.2 Adoption plan (ADKAR)

Element	For the Health Records team
Awareness	Sponsor explains why identity accuracy underpins safe AI use
Desire	Team helps design the steward queue; stewardship recognised as a skilled role
Knowledge	Training on matching evidence, decision authority, reason codes and escalation
Ability	Sandbox practice, super-user coaching, realistic workload during clean-up
Reinforcement	Queue metrics shared with the team; monthly feedback into HC4

Appendix E. Glossary

Term	Meaning
BRD	Business Requirements Document
HC1 to HC4	Human-in-the-loop checkpoints defined in Section 4
MRN	Medical record number (site-specific patient identifier)
RACI	Responsible, Accountable, Consulted, Informed
ROM	Rough order of magnitude (a preliminary cost estimate)
Shadow mode	Tool produces scores that staff can see for evaluation but no one acts on them
Tier	High, Medium or Low readmission-risk category assigned from governance-approved thresholds

Appendix F. Tools and AI assistance

This submission was developed with the assistance of Claude (Anthropic), an AI assistant, which supported drafting of documentation, diagram generation, and prototype code. All problem framing, requirements decisions, process design, and business judgement reflect the work of the Pod Nova team.

Citation (APA 7th): Anthropic. (2026). *Claude* (Sonnet 5) [Large language model]. https://claude.ai