MindBridge Whitepaper

Patient intake that feels like care

Version 1.2

Audience: Clinicians, Clinical Partners, and Investors

Region focus: Australia-first pilot pathway

December 2025

Important notices (read first):

- **Clinical decision support boundaries:** MindBridge is designed to support intake, documentation, and clinician decision-making. It is **not** a diagnostic tool and does not replace clinical judgment.
- **Safety posture:** MindBridge is designed so that risk/urgency **review cues** route to clinician review. This document describes *intended* controls and evaluation plans. Clinical impact claims require prospective validation.
- **Privacy and security:** This document describes a security posture and control intent. Compliance is determined by deployed configuration, vendor controls, and partner governance.
- **Forward-looking statements:** Capabilities described as "planned" or "intended" are not guarantees and depend on resourcing, governance, and integration scope.

Distribution: This whitepaper is intended for diligence review by potential investors and clinical partners.
Scope: Operational workflow outcomes are prioritised over unvalidated clinical effectiveness claims.

Contents

1 Executive Summary

1.1 Problem

Mental health services face sustained demand while clinician time remains scarce. Intake and triage frequently constrain throughput: referrals arrive incomplete, staff chase missing information, triage occurs with weak signals, and clinicians reconstruct histories under time pressure. The result is slower routing into care and avoidable administrative load.

1.2 Solution

MindBridge is an **AI-assisted intake and triage workflow platform** designed to:

1. capture structured patient-reported information via adaptive questionnaires;
2. convert responses into clinician-ready briefs with provenance;
3. surface risk/urgency **review cues** for clinician oversight;
4. export structured outputs aligned to FHIR Questionnaire and QuestionnaireResponse;
5. provide role-based dashboards, audit trails, and routing workflows.

1.3 Claims discipline

MindBridge is positioned as **workflow infrastructure**. The intended impact is operational:

- higher intake completeness at first clinical contact,
- reduced time-to-triage decision (with clinician review),
- lower administrative follow-up for missing information,
- improved clinician readiness for first appointment,
- measurable acceptability for patients and trust for clinicians.

MindBridge does **not** claim autonomous triage, diagnosis, or clinical outcome improvement without prospective validation.

1.4 Partnership intent

MindBridge seeks:

- **clinical advisory partnerships** to define safety boundaries, workflows, and evaluation design;
- **pilot sites** to measure operational outcomes under governance;
- **investment** to fund product hardening, integrations, compliance readiness, and multi-site validation.

2 The Intake Bottleneck: What Breaks in Real Services

2.1 Observed failure pattern (system-level)

- **Low-signal referrals:** inconsistent detail, mixed formats, unclear urgency.
- **Follow-up churn:** staff time spent gathering missing basics.
- **Triage uncertainty:** limited structured signal for prioritisation.
- **Duplication:** repeated storytelling across forms, calls, and first appointments.
- **Documentation overhead:** clinician time diverted into note reconstruction.

2.2 Root causes addressed

MindBridge targets four operational root causes:

1. information quality (completeness and structure),
2. routing uncertainty (triage signal),
3. duplication (repeat collection),
4. clinician documentation burden (first-visit preparation).

2.3 Measurable definition of success

MindBridge defines success through outcomes that clinical operations teams and investors can verify:

- intake completeness at triage review,
- time-to-triage decision,
- administrative follow-up rate,
- clinician preparation/documentation time,
- patient acceptability and clinician trust,
- safety performance of review cues (false negatives are treated as critical).

3 Product and Workflow

3.1 Core workflow

1. **Patient intake:** adaptive questionnaires, plain language, accessibility-first design.
2. **Synthesis:** clinician brief, structured narrative, and provenance links to inputs.
3. **Review cues:** risk/urgency cues surfaced for clinician review (no silent autopilot).
4. **Routing:** booking, waitlist, referral redirection, and escalation workflows.
5. **Structured export:** FHIR-aligned questionnaire outputs for portability.

3.2 Conceptual workflow diagram

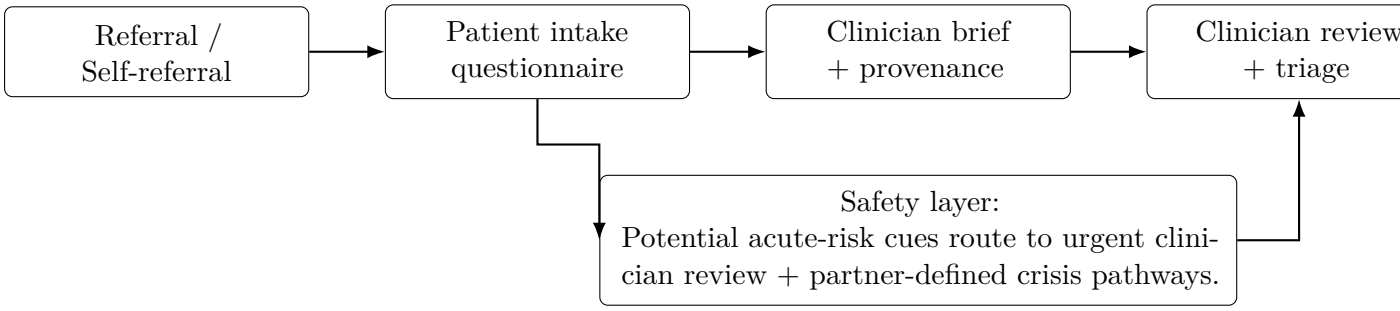

Figure 1: MindBridge intended workflow (conceptual).

3.3 Stakeholder value (diligence view)

Table 1: Value by stakeholder (and non-negotiables)

Stakeholder	Primary value	Non-negotiables
Patients	Less repetition, clearer next steps, earlier triage	Privacy, dignity, transparency, safe escalation pathways
Clinicians	Cleaner briefs, better first-visit readiness, less admin	Auditability, provenance, clinician control, no diagnosis autopilot
Operations	Reduced follow-up churn, smoother booking, clearer routing	Access control, reporting, minimal workflow disruption
Health systems	Throughput improvement, prioritisation support	Governance, interoperability, safety monitoring
Investors	Scalable workflow layer with integration upside	Claims discipline, risk controls, evaluation plan, enterprise readiness

4 Interoperability and Structured Data Capture

4.1 Interoperability objective

MindBridge is designed so intake outputs can move across systems without manual re-entry. The interoperability stance is **structured-first**.

4.2 FHIR-aligned approach (high level)

- intake instruments represented as FHIR `Questionnaire`,
- patient responses represented as `QuestionnaireResponse`,
- exports scoped by role, consent, and partner governance,
- provenance maintained between raw inputs and derived summaries.

Table 2: Example intake fields to FHIR mapping (illustrative)

Intake element	**FHIR resource**	**Notes**
Presenting concerns	QuestionnaireResponse.item	Structured selections + free text; stored with provenance
Functional impact	QuestionnaireResponse.item	Work/school, relationships, sleep, self-care
Risk screening items	QuestionnaireResponse.item	Review cues; not diagnostic claims
Treatment history	QuestionnaireResponse.item	Patient-reported history and medications
Preferences & access	QuestionnaireResponse.item	Telehealth preference, accessibility needs
Consent acknowledgements	QuestionnaireResponse.item	Explicit consent items stored for audit

5 Architecture (Diligence View)

5.1 Design goals

- **Privacy by design:** least privilege, strong tenant isolation.
- **Auditability:** derived outputs trace to source inputs.
- **Safety containment:** AI components cannot take irreversible actions.
- **Integration readiness:** structured exports (FHIR-aligned).
- **Operational resilience:** monitoring, incident response, and change control.

5.2 Logical architecture

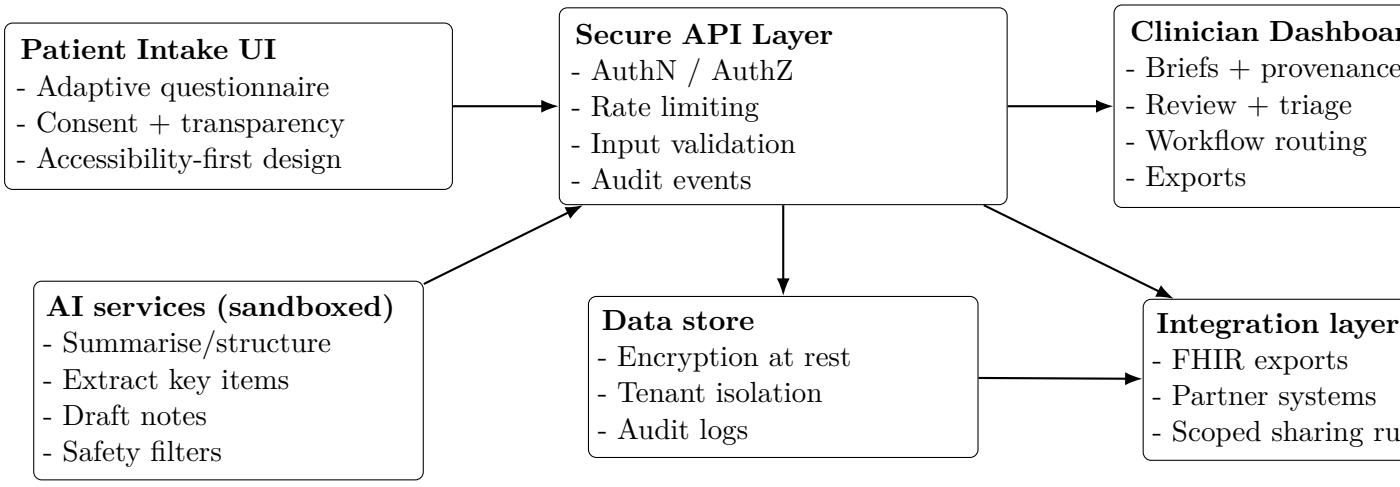

Figure 2: AI is mediated by governance and review; no safety-critical autonomy.

5.3 Provenance as a product requirement

A credible clinical workflow requires that:

- every summary statement can be traced to specific patient inputs,
- derived items are clearly labelled as derived,
- clinicians can override and annotate outputs, with audit capture.

6 Data Governance and Privacy

6.1 Governance artefacts expected by partners

Enterprise and public health partners typically require:

- data inventory (what is collected, why, where stored, retention),
- data classification and access policy,
- purpose limitation and data minimisation posture,
- sharing and export governance (who can export what, when, and why),
- audit and accountability model (event logging and retention).

6.2 Data minimisation model

MindBridge is designed around two constraints:

1. collect only data that can be justified for intake/triage/workflow support,
2. progressively disclose sensitive items only when relevant.

Table 3: Minimum data dictionary (diligence baseline)

Category	**Examples**	**Sensitivity**
Identity & contact	Name, DOB, contact method	High
Presenting concerns	Free text + structured selections	Highest
Risk screening items	Urgent review cues and safety-related items	Highest
Clinical history	Patient-reported history, medications, prior treatment	Highest
Preferences & access	Telehealth preference, accessibility requirements	High
Operational metadata	timestamps, completion state, routing state	Medium
Audit logs	access, exports, reviewer actions, configuration changes	High

6.3 Retention and deletion posture (non-trivial reality)

Clinical systems require retention policies that balance:

- clinical documentation needs,
- partner governance rules,
- audit accountability,
- lawful deletion pathways where applicable.

MindBridge is designed to support policy-driven retention and controlled deletion workflows, with audit preservation where required.

7 Security and Compliance Readiness

7.1 Threat model

Mental health intake systems are exposed to predictable threat patterns:

- credential theft and account takeover,
- tenant isolation failures,
- insecure exports (uncontrolled PDFs and broad sharing),
- insider access without auditability,
- prompt injection through free text fields,
- vendor and supply-chain risk.

7.2 Control families (diligence framing)

Table 4: Security controls and evidence expectations

Control family	Evidence typically expected
Access control	RBAC matrix, MFA posture, admin permissions, join/leave workflows
Audit logging	Access logs, export logs, reviewer decisions, config change logs
Data protection	Encryption in transit/at rest; secrets management approach
Secure SDLC	Review gates, dependency scanning, vulnerability management workflow
Monitoring	Uptime/latency monitoring; alerting; anomaly detection
Incident response	Runbooks, notification process, post-incident review practice
Change control	Release notes, rollback plan, model update governance
Vendor risk	Subprocessor inventory, minimum security expectations, contracts

7.3 Prompt injection and exfiltration controls

MindBridge is designed to reduce prompt-based attacks through:

- strict separation of system instructions from user content,
- schema-constrained outputs for high-risk summaries,
- sanitisation and detection of adversarial patterns,
- no model capability to directly access arbitrary external tools or secrets,
- audit capture of anomalous prompts and outputs for investigation.

8 Clinical Safety Case and Risk Management

8.1 Safety case structure

A defensible safety case contains:

1. intended use and boundaries,
2. hazard analysis (how harm could occur),
3. risk controls (prevention, detection, response),
4. residual risk acceptance and sign-off responsibilities,
5. post-deployment monitoring and incident learning loop.

Table 5: Hazards and controls (illustrative; tailored per pilot governance)

Hazard	**How harm occurs**	**Controls (intent)**
Missed urgent cue	Concerning patient statements are not surfaced	Conservative review cue design; clinician review gates; monitoring for near-misses
Overconfident summary	Derived statements presented as facts	Provenance labelling; calibrated language; clinician edit/override
Inappropriate routing	Incorrect urgency suggestion increases delay or burden	Suggestions only; mandatory review; documented rationale
Patient distress	Intake wording causes distress or disengagement	Plain-language testing; progressive disclosure; skip options
Privacy breach via exports	Documents shared without control or logging	Scoped export permissions; watermarks; export logs; policy
Bias / inequity	Differential performance across groups	Subgroup evaluation; monitoring; governance review and iteration

8.2 Escalation policy (partner-governed)

Escalation must be configured per clinical partner, including:

- cue thresholds for urgent review,
- clinical coverage expectations,
- patient-facing crisis instructions,
- fallback pathways when clinician coverage is unavailable.

9 Model Risk Management and Evaluation Protocol

9.1 Model risk management (MRM)

MindBridge is designed to operate under an MRM posture:

- model inventory (versions, use cases, data access boundaries),
- model cards (intended use, limits, failure modes, evaluation results),
- change control (test, approval, rollback),
- monitoring (drift, error sampling, subgroup checks).

9.2 Evaluation layers (measurable without exaggeration)

Layer 1: Output quality (offline)

- clinician-rated usefulness of briefs (rubric-based),
- factual consistency vs patient inputs (provenance checks),
- structure quality and readability.

Layer 2: Workflow outcomes (pilot)

- time-to-triage decision,
- admin follow-up rate,
- clinician prep/documentation time,
- intake completion rate.

Layer 3: Safety and reliability

- false negative rate for urgent review cues (critical),
- false positive rate (noise burden),
- near-miss reporting and root cause analysis.

Table 6: Clinician evaluation rubric (example template)

Dimension	Rating guidance (1–5)
Usefulness	Demonstrated time savings and improved readiness for first contact
Accuracy vs inputs	Claims traceable to patient responses; no hallucinated facts
Structure	Clear sections: concerns, history, function, goals, risks, next steps
Safety wording	No diagnostic language; no overconfidence; correct escalation tone
Noise level	Low irrelevant content; key items are surfaced

9.3 Bias and subgroup evaluation

Minimum diligence requires subgroup performance checks across cohorts relevant to the deployment context (subject to ethical collection and governance).

10 Implementation and Operations

10.1 Pilot implementation pathway (90-day practical)

1. **Weeks 1–2:** scope definition, governance workshop, baseline measurement plan.
2. **Weeks 3–4:** intake instrument configuration, routing workflow mapping, role setup.
3. **Weeks 5–6:** clinician training, dry runs, escalation drills.
4. **Weeks 7–10:** live pilot, weekly review meetings, quality sampling.
5. **Weeks 11–12:** outcome analysis, safety case update, rollout recommendation.

10.2 RACI (responsibility model template)

Table 7: RACI matrix (template; populated per pilot site)

Activity	**Clinic Sponsor**	**Clinical Lead**	**Privacy/Security**	**MindBridge**
Define intended use	A	R	C	R
Configure escalation policy	C	A/R	C	R
Approve intake instrument	C	A/R	C	R
User access / RBAC	C	C	A/R	R
Incident response	C	C	A/R	R
Model updates	C	C	C	A/R
Evaluation reporting	A	R	C	R

10.3 Operational monitoring (baseline)

- uptime and latency tracking,
- error rates (client and API),
- anomalous access monitoring,
- export frequency and audit sampling,
- model output sampling + clinician feedback loop.

11 Economics and ROI (Buyer Justification Framework)

11.1 ROI model (variable-based)

The most defensible ROI framing in early deployments is labour and throughput, expressed as variables that can be measured during pilots.

Table 8: ROI variables (site-specific)

Variable	Definition
N	New intakes per month
T_a	Admin minutes saved per intake
T_c	Clinician minutes saved per intake (prep/documentation)
C_a	Admin cost per hour
C_c	Clinician cost per hour
P	Monthly platform cost

Estimated monthly labour value:

$$Value = N \cdot \left(\frac{T_a}{60} C_a + \frac{T_c}{60} C_c \right) - P$$

11.2 Additional value drivers

Beyond labour, operational value often includes improved routing, reduced referral churn from incomplete information, and better first-appointment preparedness.

12 Procurement Readiness (Enterprise Reality)

12.1 Procurement artefacts commonly required

A standard procurement pack typically includes:

- security overview (architecture, controls, data flow),
- data inventory and retention posture,
- incident response plan and responsibilities,
- access control model (RBAC, MFA, audit logging),
- vendor/subprocessor inventory and processing locations,
- evaluation plan (metrics, sampling, safety monitoring),
- change management process (including model updates).

Table 9: Enterprise diligence checklist (questions typically asked)

Question	**Evidence expected**
Where is data stored and processed?	Deployment regions, vendor inventory, subprocessors
How is tenant isolation enforced?	Technical mechanism, tests, access controls
How is data protected?	Encryption posture, key management approach
Is MFA/SSO supported?	Current posture and roadmap with governance
What audit logs exist?	Access logs, export logs, reviewer and config logs
How are incidents handled?	Runbooks and notification process
How are model changes validated?	Offline evaluation, sign-off pathway, rollback plan
How are retention and deletion managed?	Policy-driven workflows and audit preservation

13 Roadmap (24-month Diligence View)

Table 10: Roadmap (indicative; dependent on governance and integration scope)

Phase	**Focus**
0–3 months	Pilot hardening: provenance UI, audit completeness, escalation configuration, instrumentation
3–6 months	Advisory loop: rubric evaluation, safety review cadence, workflow optimisation
6–12 months	Integrations: expanded FHIR exports, partner routing workflows, reporting
12–24 months	Multi-site validation, enterprise governance (SSO, policy packs), procurement readiness

14 Key Risks and Mitigations

Table 11: Risks and mitigations (investor and clinical diligence view)

Risk	**Failure mode**	**Mitigation posture**
Safety-critical errors	Missed urgent cues or misleading summaries	Review gates, conservative cues, monitoring, incident response
Clinician distrust	Perception of black box or liability shift	Provenance, edit/override, auditability, clear boundaries
Privacy breach	Unauthorised access or insecure exports	RBAC, tenant isolation, encryption, logging, testing
Workflow mismatch	Added steps instead of reduced friction	Co-design with sites, pilot iteration, measured outcomes
Unvalidated claims	Credibility loss during diligence	Claims discipline, evaluation plan, transparent reporting
Integration delays	EHR variability slows deployment	FHIR-first staging, narrow pilot scope, phased integration

15 Partnership and Investment Thesis

15.1 Why this category can win

Intake is upstream of almost every clinical workflow. Improvements at the intake layer can compound across:

- reduced administrative churn,
- improved triage signal,
- better allocation of scarce clinician time,
- stronger continuity of care through structured data.

15.2 Why MindBridge is positioned to win

Defensibility is framed around:

- safety-case design (review-gated cues, conservative posture),
- provenance and auditability (trust-building mechanics),
- interoperability-first structured capture (FHIR-aligned),
- evaluation discipline aligned to clinical governance.

A Appendix A: Example Intake Domains (Structure)

Actual question sets are configured per partner governance and cohort.

1. Identity and contact (minimal required)
2. Presenting concerns (free text + structured selections)
3. Symptoms and severity (structured)
4. Duration and trajectory
5. Functional impact
6. Risk screening cues (partner-defined; review gated)
7. Treatment history (patient-reported)
8. Preferences and access needs
9. Goals (patient-defined)

B Appendix B: Glossary

- **FHIR:** Fast Healthcare Interoperability Resources.
- **RBAC:** Role-Based Access Control.
- **MRM:** Model Risk Management.
- **Review cue:** Signal prompting clinician attention (not a diagnosis).