

AetherHeal Clinical Safety Document

From an AI proposal to a physician's review

AetherHeal Global Inc.

Issued: September 13, 2026 · v1.0

The full text behind the web summary at aetherheal.com/en/clinical-safety. The tables and detailed principles removed from the screen live here.

Purpose-built AI engines for each job in care, and a Clinical OS that keeps them on one patient state under one line of physician authority. It runs beside the chart. The existing hospital chart remains the system of record. From the first proposal to the evidence that the work finished, there is a place where the physician checks.

A design document covering review responsibilities, when a draft becomes a record, and who controls data and version changes.

1 Abstract

When the power fails you can bring it back. A wrong clinical judgment stays with the patient. That is why keeping the system non-autonomous cannot live in an ethics statement. It has to be something a physician can check and stop inside the product itself. The Golden Clinical Loop is one turn of work: the AI proposes, the physician confirms, the work runs, and evidence that it finished is written down.

We do not claim autonomous clinical decisions, guaranteed accuracy, prevention of medical disputes, completed validation, regulatory clearance, or improved clinical outcomes. The product screens shown here are synthetic data from a demo account. They are not evidence of clinical validity or of results in the field.

Document type	Clinical safety principles
Scope	Clinical OS and the Golden Clinical Loop
Author	AetherHeal

Company description

AetherHeal is building clinical execution infrastructure for hospitals and outpatient networks. Patient-access engines handle approved information, multilingual conversation, booking, and intake. Clinical-intelligence engines prepare encounter context, specialty inputs, and documentation for physician review. Care-execution engines coordinate permitted administrative actions, verification, ownership, and follow-through. These are not independent chatbots and Clinical OS is not a replacement central brain. Intelligence is distributed across bounded tasks; state, permissions, evidence, and accountability remain coordinated through a shared operating layer and a governed execution contract called the Golden Clinical Loop.

Physicians retain clinical judgment, treatment authority, correction, and final clinical responsibility. The existing hospital chart remains the system of record.

2 Scope and the responsibility boundary

Safety is not judged one answer at a time. What we review is the record that runs, in the patient's context, until the job is finished.

In clinical work the quality of a proposal does not explain safety on its own. We look at which inputs and which uncertainty the proposal came from, who reviewed it, and under whose authority it ran. We also look at what evidence of completion was left in the hospital's system of record.

Clinical OS holds patient identifier, evidence, authority, owner, and completion state together in one flow. The chart records what finished. We handle what has not finished yet. The existing hospital chart remains the system of record.

3 Execution authority × clinical impact

Execution authority and clinical impact are classified apart. One axis separates what the system may read, what it may propose, and what it may change. The other separates how far a job reaches into clinical judgment. Crossing the two sets the controls a job needs.

Table 1: Execution authority axis.

Level (code)	Definition
Read (A_READ)	Reads only approved information and changes no state in a system of record.
Propose (B_PROPOSE)	Proposes a draft or a next action for review. It does not execute.
Governed write (C_GOVERNED_WRITE)	Changes state only behind a named person's approval gate and with an allowed tool.

Table 2: Clinical impact axis.

Level (code)	Definition
Non-clinical operations (N_NONCLINICAL)	Institutional operating work that does not bear directly on clinical judgment.
Protocol-bound (P_PROTOCOL_BOUND)	Handled only under an institution-approved protocol and a named clinical reviewer.
Regulated or clinical-decision impact (R_REGULATED_OR_DIAGNOSTIC)	Kept apart as work that needs regulatory review, or the judgment of a licensed physician.

Table 3: Execution authority by clinical impact (rows = execution authority, columns = clinical impact).

Execution authority	Non-clinical operations (N_NONCLINICAL)	Protocol-bound (P_PROTOCOL_BOUND)	Regulated or clinical-decision impact (R_REGULATED_OR_DIAGNOSTIC)
Read (A_READ)	Read × Non-clinical operations (A_READ × N_NONCLINICAL)	Read × Protocol-bound (A_READ × P_PROTOCOL_BOUND)	Read × Regulated or clinical-decision impact (A_READ × R_REGULATED_OR_DIAGNOSTIC)
Propose (B_PROPOSE)	Propose × Non-clinical operations (B_PROPOSE × N_NONCLINICAL)	Propose × Protocol-bound (B_PROPOSE × P_PROTOCOL_BOUND)	Propose × Regulated or clinical-decision impact (B_PROPOSE × R_REGULATED_OR_DIAGNOSTIC)
Governed write (C_GOVERNED_WRITE)	Governed write × Non-clinical operations (C_GOVERNED_WRITE × N_NONCLINICAL)	Governed write × Protocol-bound (C_GOVERNED_WRITE × P_PROTOCOL_BOUND)	Governed write × Regulated or clinical-decision impact (C_GOVERNED_WRITE × R_REGULATED_OR_DIAGNOSTIC)

This classification exists to set the boundary between internal execution control and clinical impact. It is not automatically the same as a legal or regulatory classification.

4 Protocol-bound (P_PROTOCOL_BOUND)

Protocol-bound work needs an approved procedure and a human review. Work classified as protocol-bound enters the controlled scope only when all seven conditions below hold.

1. **Institution-approved protocol** · An institution-approved protocol is the reference.
2. **Named clinical reviewer** · A named clinical reviewer is assigned.
3. **Version** · The protocol version in force is identified.
4. **Expiry** · There is an expiry point at which the work is reviewed or stopped.
5. **Rollback** · A path back to the previously approved version stays open.
6. **Controlled rollout** · Rollout happens after the approved scope and the target institution are confirmed.
7. **No automatic updates** · A new version is never applied automatically.

5 Private build line (AGENT RAIL / COMPILER)

The compiler works only with approved tools and authority that already exists. Agent Rail / Compiler is not a publicly available product. It is a private prototype that checks tools, authority, and versions before an execution path is built.

Private prototype · We record it as one private build line and nothing more. It does not imply public availability.

1. **No arbitrary code execution** · It does not run arbitrary code.
2. **Approved tool registry only** · It references only the tools in the approved tool registry.

3. **No new authority** · Compiling never creates new authority.
4. **Static checks** · Structure and authority boundaries are checked statically before rollout.
5. **Human review** · A named person reviews the output.
6. **Versioning and rollback** · Every output is versioned, and a path back to the previous version stays open.

6 Responsibility and owner

Responsibility is split into seven layers, each with a name on it. Clinical judgment, institutional protocol, system control, and rollout approval do not collect in one place. Each responsibility has a different owner.

Table 4: Responsibility and owner.

Responsibility	Owner
Clinical judgment and treatment	The licensed clinician
Institutional protocol	The institution and its clinical reviewer
Access, tools, and the audit log	AetherHeal
Staff work	The role owner inside the institution
Source data	The owner of the source data
Review before publication	AetherHeal and the reviewer
Rollout approval	A named person's approval gate

7 Data boundary

Data use starts from the institution's boundary and a path to withdraw. Data is bound to its processing purpose, to the institution's boundary, and to a retention period. What we learn from corrections is kept apart from the boundary that forbids reusing source health information.

1. **Per-institution boundary** · Data and execution context stay inside each institution's boundary.
2. **No cross-institution reuse of source health information** · Source personal health information is never reused for another customer.
3. **Stated processing purpose and processor role** · The processing purpose and the processor role are written down.
4. **Learning from corrections and workflow patterns** · Only permitted corrections and de-identified workflow patterns are used for learning.
5. **Consent and withdrawal** · Consent status and withdrawal are recorded and applied to later processing.
6. **Retention and deletion** · Retention period, deletion conditions, and the deletion procedure are set together.

8 The problem-oriented medical record (POMR) draft and its implementation boundary

A medical record is what a later review uses to judge whether care was appropriate. A record draft exposes the structure of each problem and the gaps in it. Before it becomes the final record, a physician has to review it.

DockieTalkie Clinical Copilot lays out the record draft in problem-oriented medical record (POMR) structure. For each problem it keeps the evidence, the assessment, the plan, and the gaps apart, so the physician can see what needs checking.

1. Problem-oriented medical record (POMR) structure, with evidence, assessment, and plan kept apart for each problem
2. A check of the statutory entries: chief complaint, diagnosis, course, treatment, and date and time

We are reviewing the implementation scope for electronic signature and record-integrity safeguards against the electronic medical record requirements in Article 23 of the Medical Service Act [4]. A draft becomes a record only after a physician reviews and approves it. Fully meeting the legal electronic signature requirement, and the retention and integrity requirements for electronic medical records, is a separate implementation scope.

9 The stages of the Golden Clinical Loop and their boundaries

The Golden Clinical Loop is the one turn of work every product follows. Each stage has a condition that has to hold before the work moves on.

1. **Observe** · Source, time of observation, and uncertainty are preserved.
2. **Structure** · A statement is never promoted to a confirmed fact on its own.
3. **Propose** · A proposal is a scoped next action, or a point to review. Nothing more.
4. **Confirm** · Authority is granted explicitly and its scope is set. The default is off.
5. **Execute** · Only the approved scope runs, inside an allowlist and idempotency controls.
6. **Verify** · We do not take the system's word that something finished. We check the processing record or the state change in the hospital's system of record: booking, EMR, CRM.
7. **Record** · The execution record and any corrections are written down.
8. **Follow up** · Owner, due date, and the escalation path for exceptions are kept. Whatever is left carries into the next review and into the configuration.

10 Authority by stage

The Golden Clinical Loop is folded into five stages so the authority of the AI, the physician, staff, and the system can be told apart. Proposal and execution continue only inside the scope the physician set. Completion requires evidence from the hospital's system of record.

Table 5: Authority by stage.

Stage	AI	Physician	Staff	System
Observe and structure	Preserves source, time of observation, and uncertainty, and does not promote a statement to a confirmed fact on its own.	Reviews the evidence, assessment, plan, and gaps in the record draft.	Carries out the checks assigned to them.	Holds the patient identifier and the evidence in one flow.
Reason and propose	Proposes only a scoped next action or a point to review.	Approves, revises, or rejects the proposal, and makes the final clinical judgment.	Confirms the assigned task and who owns it.	Keeps the proposal's inputs, its uncertainty, and the record of review.
Confirm	Grants no authority, and holds none to execute directly.	Grants authority explicitly and sets its scope.	Confirms the assigned task and the scope of authority.	Defaults to off and holds only the authority that was granted explicitly.
Execute	Holds no authority to execute directly.	Confirms the approved scope of execution, and revises or rejects it where needed.	Carries out the assigned task.	Runs only the approved scope, inside an allowlist and idempotency controls.
Complete and track	Cannot declare that anything is finished.	Reviews exceptions and corrections, and makes the final clinical judgment.	Works the follow-up by owner and due date, and escalates exceptions.	Marks a job complete only when there is a processing record or a state change in the hospital's system of record.

11 Claim scope and the evidence it requires

Raising the level of a claim requires new evidence and a new review. Nothing we say about a product goes past the evidence we hold today.

Table 6: Claim scope and required evidence.

Level	Claim scope	Evidence required
1. Workflow support	Interpretation, transcription, summary, glossary, report drafts	Reliability, usability, error records, and physician review
2. Decision support for physicians	Questions, risk signals, checklists, structured references	Human factors evaluation, records of physician action, and safety review
3. Validated vertical use	Indication-specific support inside a defined workflow	Performance evaluation, subgroup analysis, and a clinical protocol
4. Regulated clinical claims	Claims about identifying disease or about the effect on treatment	Clearance, a quality system, clinical evidence, and post-market monitoring

12 The boundary of what we say in public

Public wording carries both the state of the product and the level of evidence behind it. The lists below separate what we can say today from what we do not say.

What we can say

- Physician-supervised clinical decision support
- Support for record completeness and risk management
- Support for reference, workflow, classification, and measurement
- An evidence generation plan and regulatory review

What we do not say

- The AI makes a patient's clinical decision on its own
- Accuracy is guaranteed
- It prevents medical disputes
- Clinical validation or regulatory clearance is complete
- It improves clinical outcomes
- It replaces every EMR immediately

The external standards to read alongside this section are the WHO guidance on the ethics and governance of AI for health [1], the FDA guidance on clinical decision support software [2], and the IMDRF document on clinical evaluation of software as a medical device [3].

13 References

1. World Health Organization. Ethics and governance of artificial intelligence for health: WHO guidance. <https://www.who.int/publications/i/item/9789240029200>
2. U.S. Food and Drug Administration. 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>
3. International Medical Device Regulators Forum. Software as a Medical Device: Clinical Evaluation. <https://www.imdrf.org/documents/software-medical-device-samd-clinical-evaluation>
4. Republic of Korea. Medical Service Act, Article 23 (Electronic medical records). 의료법 제23조(전자의무기록). Korean Law Information Center.

Questions about the scope of this document or about its citations go to drjeeju@aetherheal.com.

Physicians retain clinical judgment, treatment authority, correction, and final clinical responsibility. The existing hospital chart remains the system of record.