AetherHeal: The Vertical AI Platform for Medicine

AetherHeal Global Inc.
Business Registration No.: 461-81-04200
Address: 865, Bangeojinsunhwando-ro, Dong-gu, Ulsan,
Republic of Korea
Representative: Jee Hoon Ju, M.D.

June 28, 2026

Abstract

AetherHeal is a staged, physician-supervised vertical AI platform for medicine. The narrow waist is Dockie-talkie, the reusable clinical port where conversation, terminology, translation logs, reports, and future reasoning artifacts enter the platform. Clinical Copilot supports physician-supervised reasoning over that port, while Clinical OS receives reviewed artifacts into records and operations. DermatoScan AI and WoundScan AI are validation-bounded proof points for specialty depth. This paper sets out the problem, architecture, product status, validation discipline, go-to-market strategy, and evidence-generation plan.

1 Introduction

Medical AI systems are increasingly capable of producing plausible text, classifying images, and summarizing clinical context. This is not sufficient for deployment in actual care. A clinical product must preserve provenance, fit workflow, support physician judgment, record what was accepted or rejected, and maintain a defensible boundary between assistance and autonomous medical decision-making. The core thesis of AetherHeal is that medical AI value accrues to systems that own the clinical loop, not merely the model output.

In this paper, a *clinical loop* means the full chain through which a patient explains a problem, a clinician probes and clarifies, context is refined, risk is assessed, findings are documented, consent is explained, follow-up is scheduled, and liability-sensitive records are created. The platform thesis is to make this loop structured, auditable, and physician-supervised. A system that can observe only one point in the loop is useful but not defensible. A system that can connect conversation, image, physician action, record, and outcome is qualitatively different.

AetherHeal is neither an abstract AI-model company nor a generic clinical operating system. It is a vertical AI platform with one narrow waist: Dockie-talkie creates the standard clinical port through which real encounters become structured artifacts. Clinical Copilot reasons over those artifacts for physicians, and Clinical OS connects accepted artifacts to records and operations. Specialty products extend that loop into workflows where images, devices, and clinical context matter.

2 Core Thesis

AetherHeal is the Vertical AI Platform for Medicine: a platform where clinical conversations, images, biosignals, explanations, consents, AI suggestions, physician actions, and record-and-operations data become one auditable clinical loop. The thesis can be summarized in five claims.

1. **Healthcare is uncertainty management.** The platform helps clinicians ask better questions, notice weak signals, and document reasoning before error compounds.

2. **Trust is the product.** Each suggestion must preserve context, provenance, uncertainty, and the physician action that accepted, edited, or rejected it.
3. **Workflow beats model novelty.** A clinically embedded model with adoption and audit trails can outperform a stronger generic model outside the encounter.
4. **Vertical data becomes the moat.** Specialty terminology, device context, correction logs, and outcomes are hard to copy.
5. **Non-autonomous is strategic.** Physician-supervised decision support is safer, more credible, and more adoptable than replacement claims.

The updated source narrative tightens the category language. External material presents Dockie-talkie as the narrow waist and first standard clinical port. Clinical OS is the record-and-operations layer that receives structured data from that port.

This distinction matters because platform companies often fail by attempting to sell the whole operating layer before a clinic has a narrow and urgent reason to adopt. AetherHeal reverses the sequence: start at a concrete workflow pain, capture structured artifacts there, then expand into reasoning, documentation, operations, and specialty verticals. That order is the strategy, commercial and technical alike.

3 Problem

The initial wedge is not medical tourism as a company identity. It is the operational gap faced by clinics that want foreign-patient capacity but cannot yet justify full-time interpreters or specialty-trained multilingual staff. Official Korean government reporting indicates that foreign patients visiting Korea surpassed 2 million in 2025, with 2.01 million foreign patients and 2.72 million total patient visits [Ministry of Health and Welfare, Republic of Korea, 2026].

The problem is not simply that patient and physician speak different languages. High-stakes medical explanations include procedure names, medication names, contraindications, consent language, adverse-effect warnings, and clinic-specific terms. If the message is summarized or distorted, the physician loses visibility into what the patient actually understood. Interpreter errors can carry clinical consequences, and professional interpreters improve care for limited-English-proficiency patients [Flores et al., 2003, Karliner et al., 2007].

A second problem is documentation and workflow burden. Time-motion and EHR-log studies show substantial physician time spent on electronic records and desk work, while newer ambient AI work evaluates documentation-burden reduction [Sinsky et al., 2016, Arndt et al., 2017, Husa et al., 2026]. In practice, documentation burden reaches beyond productivity. It affects record completeness, consent clarity, handoff quality, physician fatigue, and the ability to reconstruct what happened in a disputed encounter.

A third problem is the gap between promising AI research and clinically governed deployment. Medical AI reviews describe the opportunity for multimodal and generalist systems, but emphasize deployment complexity [Rajpurkar et al., 2022, Acosta et al., 2022, Moor et al., 2023]. Reporting guidelines such as CONSORT-AI and DECIDE-AI emphasize transparent evaluation, human factors, and clinical utility [Liu et al., 2020, Vasey et al., 2022]. AetherHeal's thesis is designed around this gap: AI outputs become useful when they are tied to workflow, physician review, evidence, and records.

4 Design Principles

The platform is governed by design principles that are deliberately conservative. First, AI outputs must be evidence-linked. A recommendation, question, red flag, or summary is traceable to the

encounter artifact that produced it. Second, AI outputs must be physician-actionable. The platform records whether a suggestion was accepted, dismissed, edited, marked as asked, or routed to follow-up. Third, the system must separate current product status from roadmap ambition. A prototype reasoning board is not a validated diagnostic device; a physician-only reference workflow is not a patient-facing regulated product.

Fourth, the system minimizes hidden infrastructure claims. External versions avoid overemphasizing internal frameworks, databases, model names, or cloud details. Those details matter operationally but do not define the category. The externally relevant claim is that structured sessions, participants, utterances, glossary terms, reports, usage, cost telemetry, physician actions, and record artifacts can be captured as part of a clinical loop. Fifth, the product progresses through an evidence ladder. Workflow assist is separated from physician-facing clinical decision support, and validated vertical use is separated from regulated clinical claims.

Table 1: *Design principles and their operational meaning.*

Principle	Operational meaning	Reason
Evidence-linked output	Suggestions are tied to transcript, image, device, or record artifacts.	Physicians need to inspect the basis of the AI output.
Physician action logging	Accept, dismiss, edit, mark asked, and follow-up actions are recorded.	Supervision becomes measurable rather than rhetorical.
Status discipline	Production, QA, prototype, scaffold, and roadmap are separated.	External claims stay aligned with code reality and validation.
External stack minimization	Internal frameworks and model names are not the public category.	The product is the clinical loop, not the implementation stack.
Evidence ladder	Claims advance only with matching validation and governance.	Medical AI adoption depends on safety, trust, and intended use.

5 Architecture

The architecture is staged around one standard clinical port. Layer 1 is the shared horizontal infrastructure: Dockie-talkie converts clinical conversation and procedures into structured clinical data, and Clinical Copilot turns that live context into real-time, non-autonomous, physician-supervised reasoning artifacts. Clinical OS is not Layer 1 and not Layer 2. It is the record-and-operations layer that receives the structured data the port produces and connects it across booking, reception, encounter, billing, and follow-up. Layer 2 contains vertical applications that plug into the port: DermatoScan AI as the flagship image vertical, WoundScan AI, then essential-care verticals in internal medicine and emergency care, ophthalmology, dentistry, and other specialties.

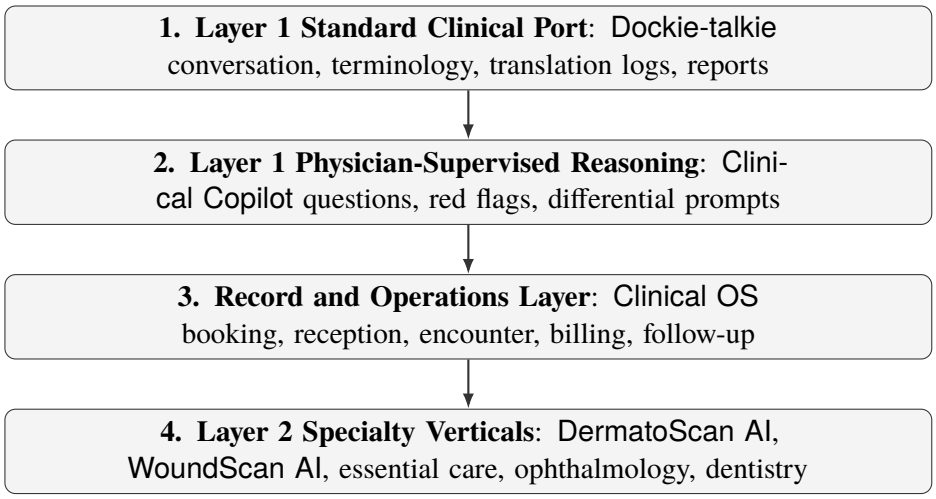

Figure 1: *AetherHeal architecture: the standard clinical port feeds physician-supervised reasoning, then the record-and-operations layer, then specialty verticals.*

A second moat is owned clinical practice. The founder is a practicing physician with owned clinical sites: two outpatient aesthetic and dermatology clinics, plus Tongyeong Seoul Hospital as an essential-care testbed for internal medicine and emergency care, selected on June 19, 2026. The pre-launch advantage is practical rather than promotional: the team can develop, apply in real practice, and quantify in-house, closing the full loop without depending on an external PoC site.

6 Product Status

The code-grounded reality is uneven by design. That unevenness is a strength when described accurately: one near-commercial wedge, one reasoning prototype, one operating foundation, and two vertical proof points moving toward validation.

The product capture in Figure 2 places Clinical OS in its proper role: the record-and-operations layer, not the company category. The interface is organized around daily clinic operations, patient flow, records, and structured clinical context.

Table 2: *Code-grounded product maturity from the updated master narrative.*

Product	Maturity	External wording discipline
Dockie-talkie	SaaS wedge	Medical interpretation SaaS with clinic-specific terminology, reports, usage limits, and a planned July 1, 2026 clinic free-opening followed by an August 2026 paid-conversion target.
Dockie-talkie Mobile	Native shell	iOS/Android shell strategy for clinician use; core features stay in the web SaaS.
Clinical Copilot	QA prototype	Physician-facing consultation copilot; current artifact is a live transcript workbench and reasoning board, not autonomous diagnosis.
Clinical OS	Operating foundation	App/portal foundation being evolved into an AI-native record-and-operations layer; vertical integrations are being consolidated.
DermatoScan AI	Reference workflow	Dermatoscope-aware clinical reference tool with target-stack rebuild in progress; external deployment requires regulatory and legal review.
WoundScan AI	P0 scaffold	Wound measurement, tissue-percentage path, and red-flag logic prototype/specification; validation and partner review are still required.

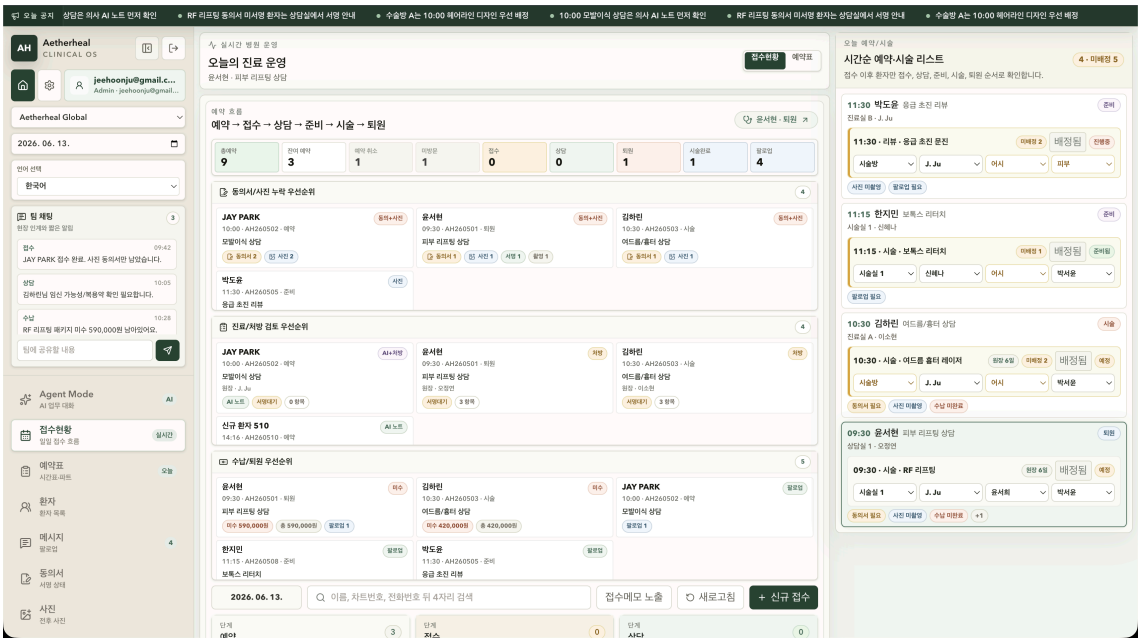

Figure 2: *Clinical OS product capture: an AI-native EMR surface for clinic operations. The intended role is to receive physician-reviewed artifacts from the clinical port and connect them to records, scheduling, and operational workflows.*

This status matrix matters because the strongest external narrative is not equal maturity across products. Dockie-talkie is the fastest route to clinic adoption and payment. Reasoning and records build on top of it through Clinical Copilot and Clinical OS, and DermatoScan AI and WoundScan AI carry the platform into image- and device-heavy specialties.

7 Clinical Workflow

Dockie-talkie is medical communication infrastructure, not a translator. Its value lies less in multilingual output than in what it preserves: the physician's actual explanation, clinic-specific terminology, reviewable transcripts and reports, and a record the operating layer can later use. It lowers the fixed-cost barrier to foreign-patient care and creates the first structured communication logs for the platform.

A representative Dockie-talkie encounter has five stages. The clinician begins a session with a patient and selects or inherits clinic-specific terminology. Speech is translated in real time while the system preserves utterances, participants, and language direction. Specialty terms are matched against the clinic glossary. The session produces a report or transcript artifact. Finally, the artifact can be reviewed, stored, and later connected to Clinical OS or Clinical Copilot. This is a narrow workflow, but it is also a repeatable clinical data port.

Figure 3 shows this wedge in product form. The capture is intentionally simple: a clinic user starts a session, language direction is preserved, utterances stay reviewable, and the resulting communication log becomes a reusable artifact.

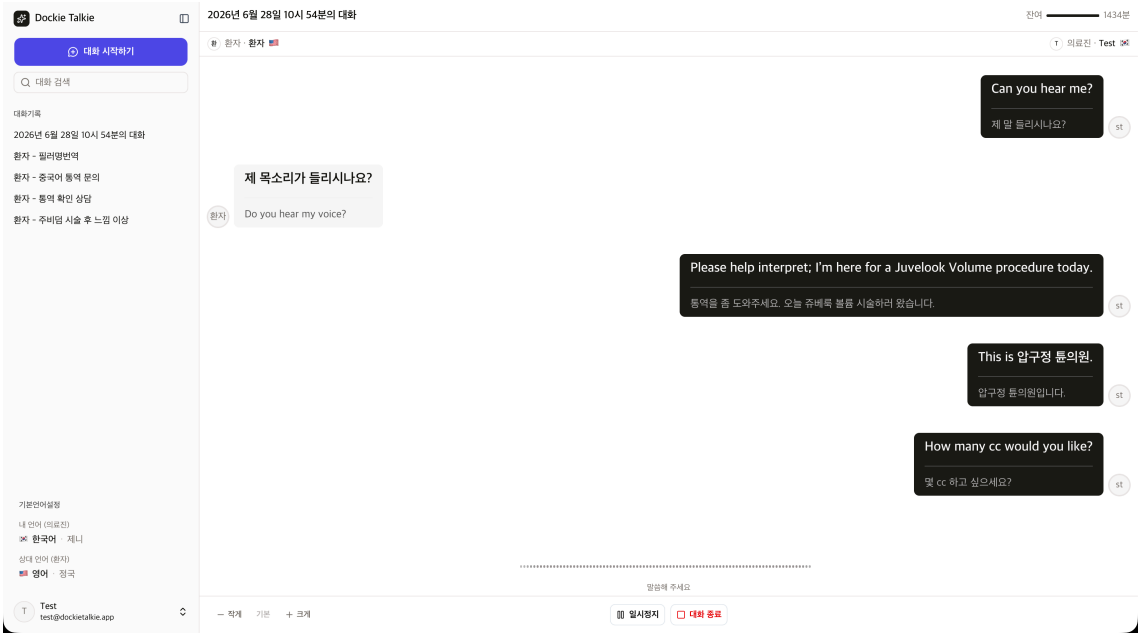

Figure 3: *Dockie-talkie product capture: a medical interpretation workflow that preserves reviewable conversation artifacts for later reasoning and documentation.*

Clinical Copilot is physician-facing and never autonomous. It turns a live encounter into questions, red flags, differential considerations, and follow-up steps; the physician accepts, dismisses, marks as asked, or revises each one. This aligns with clinical AI evaluation guidance that emphasizes workflow, human factors, and clinical utility [Vasey et al., 2022, Rajpurkar et al., 2022]. A useful copilot helps the clinician notice, check, and document the reasoning process without replacing clinical judgment.

Figure 4 presents Clinical Copilot as a physician workbench. The capture shows encounter context being turned into a differential list and follow-up prompts, preserving the non-autonomous boundary: AI produces reviewable artifacts, while the physician keeps responsibility for acceptance and final documentation.

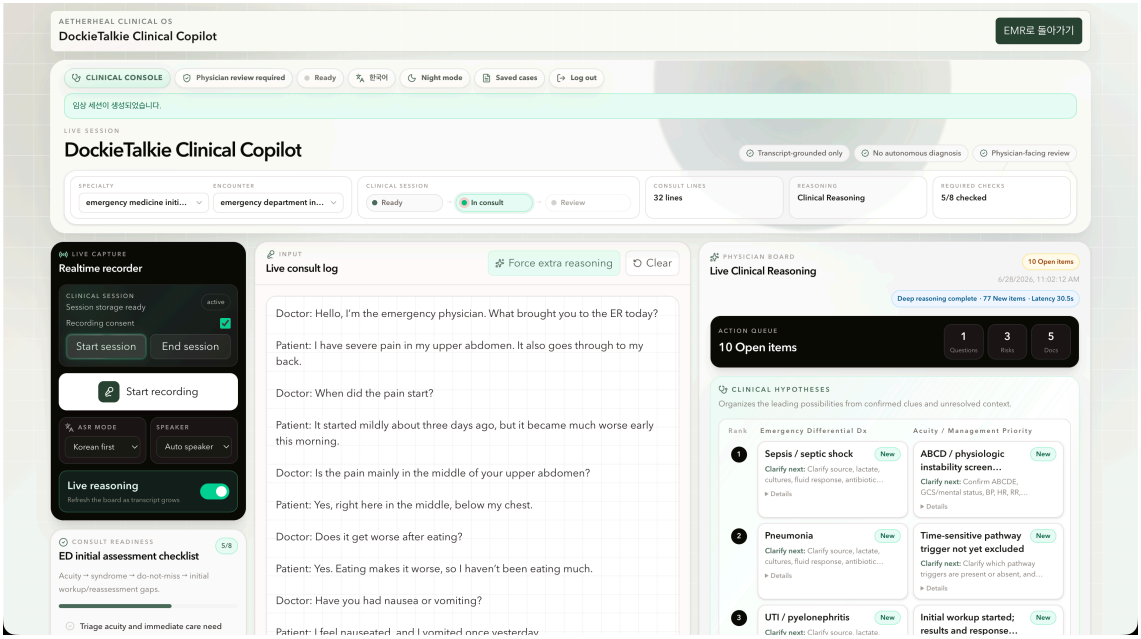

Figure 4: *Clinical Copilot product capture: differential-list generation inside a physician-facing reasoning workbench. The workflow is designed to generate reviewable reasoning and documentation artifacts rather than autonomous diagnosis.*

DermatoScan AI and WoundScan AI show that AetherHeal reaches beyond horizontal documentation. Dermatology and wound care are image-rich domains where clinical utility depends on capture conditions, device context, physician review, and outcome validation. Literature supports promise in controlled image tasks while highlighting bias, deployment, and validation gaps [Esteva et al., 2017, Laiouar-Pedari et al., 2026, Daneshjou et al., 2021, Anisuzzaman et al., 2022]. Future electronic-stethoscope integration remains a roadmap item; digital auscultation research shows promise, but product claims require indication-specific validation and regulatory review [Chorba et al., 2021].

Figure 5 shows how a specialty vertical can sit on top of the shared platform. The relevant product claim is a reference workflow: image capture, lesion context, device-linked capture conditions, physician review, and a future evidence-generation pathway connected in one place.

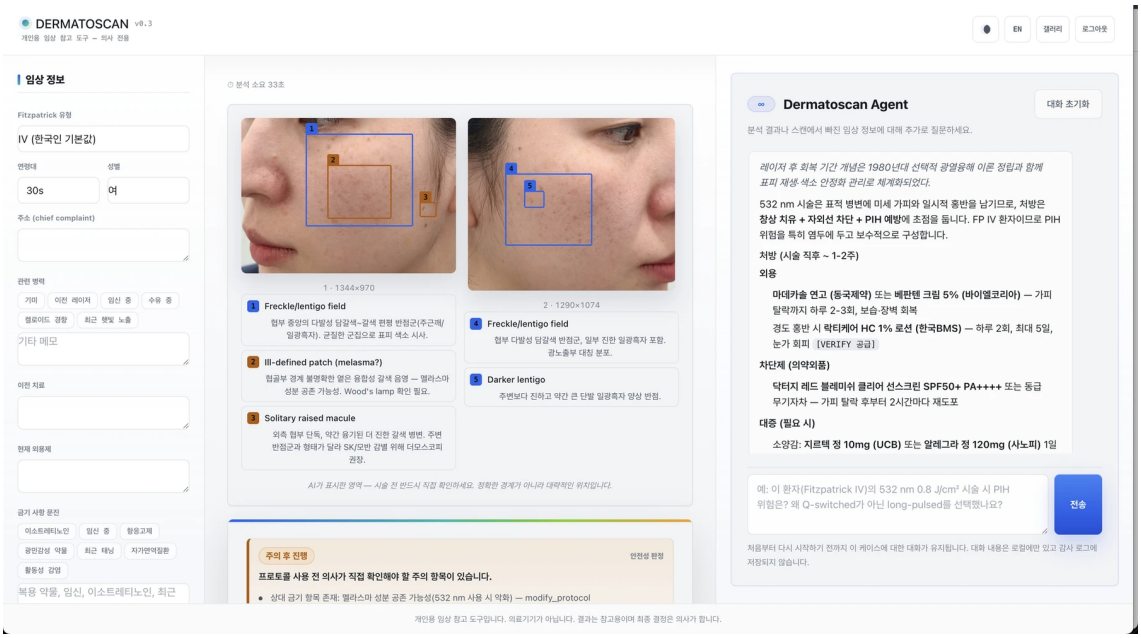

Figure 5: *DermatoScan product capture: an image-capture and agent-assisted dermatology review workflow. The capture supports the specialty-vertical thesis while remaining within explicit validation boundaries.*

8 Data Flywheel

The moat is not raw data volume alone. It is specialty-specific context tied to physician action and outcomes. Examples include glossary terms, accepted or rejected suggestions, edited notes, dermatoscope-linked cases, wound measurements, follow-up outcomes, and audit trails.

Table 3: *Clinical-loop data objects that create defensible platform value.*

Data object	Examples	Strategic value
Conversation context	utterances, speaker roles, language, translated phrases, glossary hits	Specialty terminology and explainable communication logs.
Reasoning artifacts	questions, red flags, differential prompts, checklist items	Measurable physician-supervised reasoning quality.
Physician actions	accept, dismiss, mark asked, edit, follow-up	Supervised correction signal and defensible audit trail.
Image/device context	dermatoscope magnification, wound area, tissue class, capture conditions	Vertical model improvement and quality control.
Record outputs	note, consent evidence, plan, patient instructions, referral, billing context	Integration into clinic operations and medico-legal documentation.
Outcome signals	revisit, complication, response, healing trajectory, patient-reported issue	Validation substrate for future clinical utility studies.

The key is that these objects are not independent. A translated explanation may lead to a physician action. A physician action may create a note. A note may trigger follow-up. Follow-up may produce an outcome. An outcome may refine the next suggestion. The data flywheel therefore depends on structured capture, not passive accumulation. It also depends on governance: data is de-identified,

consented, and used only under explicit compliance boundaries when it moves beyond immediate care operations.

9 Validation Strategy

AetherHeal's validation strategy is staged by intended use. For workflow-assist functions, the key evidence is operational: latency, uptime, report completion, glossary accuracy, user retention, and willingness to pay. For physician-facing reasoning functions, the key evidence is safety and usability: physician acceptance rate, dismissal rate, unsafe suggestion rate, missed red flags, and time to documentation. For image-based specialty functions, the key evidence includes capture quality, inter-rater agreement, subgroup performance, calibration drift, and follow-up outcomes.

Table 4: *Evidence ladder for product claims.*

Level	Permitted claim type	Evidence needed before moving up
Workflow assist	Translation, transcription, summarization, glossary, report draft.	Reliability, usability, error logging, and physician review workflow.
Physician-facing CDS	Questions, red flags, checklists, structured references.	Human-factors evaluation, physician action logs, safety review, and prospective pilot.
Validated vertical use	Indication-specific specialty support under defined workflow.	Performance evaluation, subgroup analysis, clinical protocol, and governance review.
Regulated clinical claim	Diagnostic or treatment-impact claim.	Regulatory pathway, quality system, clinical evidence, and post-market monitoring.

This evidence ladder prevents the company from collapsing all AI capabilities into a single overbroad medical claim. It also gives investors, government reviewers, and clinical partners a clear view of how the platform moves from workflow utility to validated specialty use.

10 Safety and Governance

The regulatory and ethical posture is strict: AetherHeal does not replace clinicians, does not autonomously diagnose, does not guarantee accuracy, and does not promise malpractice prevention. It helps clinicians structure information, review suggestions, and maintain auditable records. This posture is consistent with WHO ethical principles, IMDRF SaMD clinical-evaluation expectations, and FDA CDS guidance around clinician review and recommendation transparency [World Health Organization, 2021, International Medical Device Regulators Forum, 2017, U.S. Food and Drug Administration, 2022].

The governance layer maintains product-by-product intended-use statements. Each product carries a clear status, claims boundary, data-retention policy, physician-review requirement, and validation plan. The purpose of governance is not to slow the product down. It makes adoption possible by giving clinicians and partners confidence that the system knows the difference between workflow assistance, clinical decision support, and regulated medical claims.

Table 5: *External language guardrails.*

Use this language	Avoid this language
Physician-supervised clinical decision support	AI autonomously diagnoses patients
Supports documentation completeness and risk management	Prevents malpractice or guarantees lawsuit defense
Reference, triage-support, measurement-support, workflow-support	Fully validated diagnostic product before validation
Planned launch, paid-conversion target, pilot, QA, scaffold, target-stack rebuild	Commercialized or clinically validated unless already true
Evidence-generation roadmap and regulatory review	Guaranteed accuracy or replacement of clinicians

11 Go-to-Market

The near-term GTM is wedge-first: launch Dockie-talkie where clinics have visible pain and fast willingness-to-test, then use communication logs and clinic relationships to introduce Clinical Copilot, Clinical OS, and specialty verticals. This avoids selling a broad operating platform before a clinic has a concrete reason to adopt.

The wedge customer is a clinic that wants to serve foreign patients but is not ready to hire a full-time interpreter, or has interpreters but lacks consistent specialty terminology and reviewable communication logs. The initial value proposition is direct: start foreign-patient care with lower fixed cost, reduce interpretation variance, preserve physician explanation, and build a reusable glossary. The expansion value proposition is that these communication artifacts become the foundation for documentation, consent, follow-up, and clinical reasoning.

1. **2026: Wedge launch.** Execute the planned July 2026 free clinic opening and August 2026 paid-conversion target. Validate glossary use, interpretation quality, usage, willingness to pay, and retention.
2. **2027: Copilot and operations.** Harden the physician workbench, connect structured logs to records and follow-up, and measure physician action rates.
3. **2028: Vertical packages.** Advance DermatoScan and WoundScan through clinical partner validation and package specialty workflows.
4. **2029 onward: Platform scale.** Expand regionally, deepen data governance, validate vertical AI use cases, and add partner integrations.

12 Business Model

The business model has three layers. The first is SaaS: monthly clinic subscriptions and usage-based interpretation or reasoning capacity. The second is device and imaging: specialty packages linked to dermatoscopes, wound capture, and future biosignal workflows. The third is data and validation: governed evidence-generation partnerships using de-identified, consented, compliance-controlled clinical-loop data.

Our forecasts separate recognized revenue, pilots, planned launches, signed partners, and hypotheses. In external materials, any revenue figure that is not recognized revenue is labeled a hypothesis. This discipline is especially important because the product stack contains multiple maturity levels. Dockie-talkie may generate near-term willingness-to-pay evidence. Clinical Copilot may first generate physician-usage and safety evidence. DermatoScan AI and WoundScan AI may first generate

clinical partner validation evidence before broad commercialization.

13 Milestones

The next eighteen to twenty-four months convert narrative risk into measured progress. The core milestones are evidence-producing events, not merely feature releases. They include paid conversion, glossary quality, physician action logs, documentation time, capture quality, partner validation, and governance review.

Table 6: *Milestones that convert narrative risk into measurable evidence.*

Workstream	Milestone	Evidence produced
Dockie-talkie	Execute planned free-opening and paid-conversion target.	Usage cohorts, willingness-to-pay data, terminology error reports, retention.
Clinical Copilot	Move from workbench prototype to physician QA pilot.	Action acceptance/dismissal rates, missed-question audits, safety logs.
Clinical OS	Connect accepted encounter artifacts to records and operations.	Documentation time, record completeness, audit-trail exports.
DermatoScan AI	Complete target-stack rebuild and reference evaluation protocol.	Case library, device-context capture metrics, bias and QC checklist.
WoundScan AI	Convert P0 scaffold into partner-reviewed validation protocol.	Wound-area agreement, red-flag concordance, tissue-percentage QA.
Governance	Create intended-use register and evidence ladder.	Claims matrix, regulatory-risk memo, data-governance SOP.

14 Risks and Mitigations

The principal risk is overclaiming. AetherHeal avoids language suggesting autonomous diagnosis, guaranteed accuracy, malpractice prevention, or completed validation before evidence exists. The mitigation is a product-by-product claims matrix and a habit of separating production, QA, prototype, scaffold, and roadmap status.

A second risk is selling too broad a platform too early. Dockie-talkie mitigates that risk by proving adoption in a concrete workflow before the full record-and-operations layer is sold. A third risk is data fragmentation; the answer is standardization of the clinical artifacts that matter: utterances, glossary terms, reports, physician actions, image context, notes, follow-up, and outcomes. A fourth risk is model commoditization. AetherHeal's answer is ownership of the workflow, domain terminology, physician correction loop, and validation evidence.

15 Conclusion

AetherHeal compounds from a first clinic relationship into a staged vertical AI platform. Dockie-talkie creates the standard clinical port. From that port, physician reasoning, records, operations, and specialty verticals build around the same loop. The company wins by keeping that loop tight: real clinicians, real workflows, explicit claims, measured outcomes, and disciplined evidence generation.

16 Product Surface Gallery

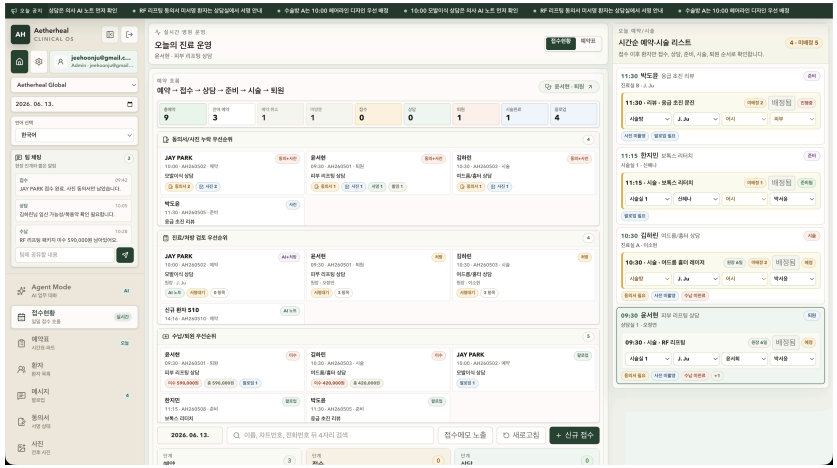

Clinical OS

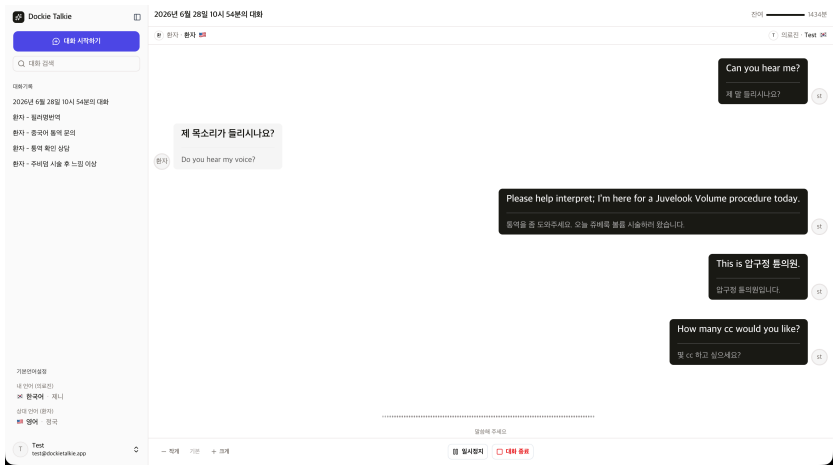

Dockie-talkie

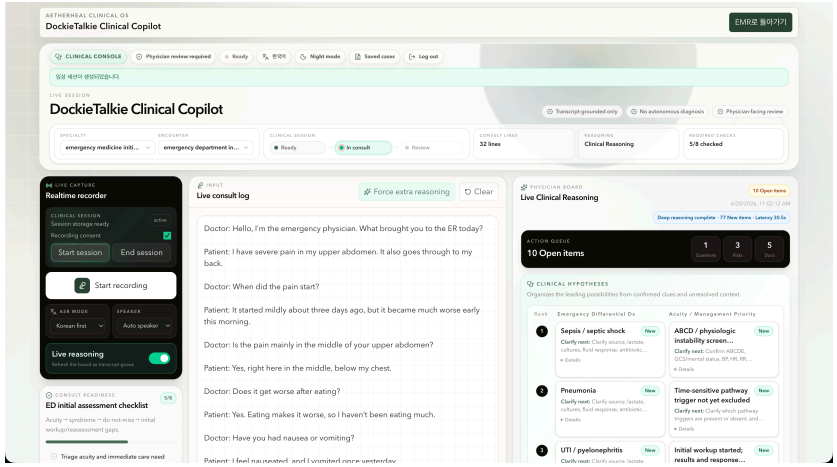

Clinical Copilot

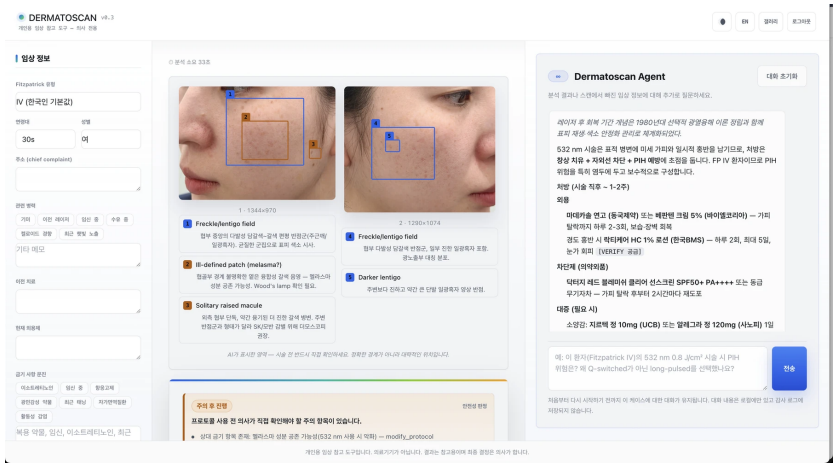

DermatoScan

Figure 6: *Compact product surface gallery using full-resolution representative screenshots rather than compressed multi-screenshot panels.*

References

Julián N. Acosta, Guido J. Falcone, Pranav Rajpurkar, and Eric J. Topol. Multimodal biomedical ai. *Nature Medicine*, 28:1773–1784, 2022. doi: 10.1038/s41591-022-01981-2.

D. M. Anisuzzaman, Chuanbo Wang, Behrouz Rostami, Sandeep Gopalakrishnan, Jeffrey Niezgoda, and Zeyun Yu. Image-based artificial intelligence in wound assessment: A systematic review. *Advances in Wound Care*, 11(12):687–709, 2022. doi: 10.1089/wound.2021.0091.

Brian G. Arndt, John W. Beasley, Michelle D. Watkinson, Jonathan L. Temte, Wen-Jan Tuan, Christine A. Sinsky, and Valerie J. Gilchrist. Tethered to the ehr: Primary care physician workload assessment using ehr event log data and time-motion observations. *Annals of Family Medicine*, 15 (5):419–426, 2017. doi: 10.1370/afm.2121.

John S. Chorba, Avi M. Shapiro, Le Le, John Maidens, John Prince, Steve Pham, Mia M. Kanzawa, Daniel N. Barbosa, Caroline Currie, Catherine Brooks, James D. Thomas, et al. Deep learning algorithm for automated cardiac murmur detection via a digital stethoscope platform. *Journal of the American Heart Association*, 10(9):e019905, 2021. doi: 10.1161/JAHA.120.019905.

Roxana Daneshjou, Mary P. Smith, Mary D. Sun, Veronica Rotemberg, and James Zou. Lack of transparency and potential bias in artificial intelligence data sets and algorithms: A scoping review. *JAMA Dermatology*, 157(11):1362–1369, 2021. doi: 10.1001/jamadermatol.2021.3129.

Andre Esteva, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542:115–118, 2017. doi: 10.1038/nature21056.

Glenn Flores, M. Barton Laws, Sandra J. Mayo, Barry Zuckerman, Milagros Abreu, Leonardo Medina, and Eric J. Hardt. Errors in medical interpretation and their potential clinical consequences in pediatric encounters. *Pediatrics*, 111(1):6–14, 2003. doi: 10.1542/peds.111.1.6.

Robyn A. Husa, John Haggerty, Andrew W. Nute, Julie Levine, Kevin Love, Xochitl Martinez, and Canada Parrish. Ambient artificial intelligence use and clinician documentation burden, productivity, and efficiency. *JAMA Network Open*, 9(5):e2615762, 2026. doi: 10.1001/jamanetworkopen.2026.15762.

International Medical Device Regulators Forum. Software as a medical device (samd): Clinical evaluation, 2017. URL `https://www.imdrf.org/documents/software-medical-device-samd-clinical-evaluation`. IMDRF/SaMD WG/N41FINAL:2017; accessed 2026-06-28.

Leah S. Karliner, Elizabeth A. Jacobs, Alice Hm Chen, and Sunita Mutha. Do professional interpreters improve clinical care for patients with limited english proficiency? a systematic review of the literature. *Health Services Research*, 42(2):727–754, 2007. doi: 10.1111/j.1475-6773.2006.00629.x.

Sara Laiouar-Pedari, Arlene Kühn, Christoph Wies, et al. Prospective evidence on artificial intelligence-assisted melanoma diagnostics: A systematic review and meta-analysis. *JAMA Dermatology*, 162(5):478–487, 2026. doi: 10.1001/jamadermatol.2026.0217.

Xiaoxuan Liu, Samantha Cruz Rivera, David Moher, Melanie J. Calvert, Alastair K. Denniston, and SPIRIT-AI and CONSORT-AI Working Group. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: The consort-ai extension. *Nature Medicine*, 26: 1364–1374, 2020. doi: 10.1038/s41591-020-1034-x.

Ministry of Health and Welfare, Republic of Korea. Foreign patients surpass 2 million in 2025, positioning korea as asia's medical tourism hub, April 2026. URL https://www.korea.net/Government/Briefing-Room/Press-Releases/view?articleId=2016788&insttCode=A260111&type=N. Korea.net press release; accessed 2026-06-28.

Michael Moor, Oishi Banerjee, Zahra S. H. Abad, Harlan M. Krumholz, Jure Leskovec, Eric J. Topol, and Pranav Rajpurkar. Foundation models for generalist medical artificial intelligence. *Nature*, 616:259–265, 2023. doi: 10.1038/s41586-023-05881-4.

Pranav Rajpurkar, Emma Chen, Oishi Banerjee, and Eric J. Topol. Ai in health and medicine. *Nature Medicine*, 28:31–38, 2022. doi: 10.1038/s41591-021-01614-0.

Christine Sinsky, Lacey Colligan, Ling Li, Mirela Prgomet, Sam Reynolds, Lindsey Goeders, Johanna Westbrook, Michael Tutty, and George Blike. Allocation of physician time in ambulatory practice: A time and motion study in 4 specialties. *Annals of Internal Medicine*, 165(11):753–760, 2016. doi: 10.7326/M16-0961.

U.S. Food and Drug Administration. Clinical decision support software: Final guidance, 2022. URL https://www.fda.gov/media/162880/download. Accessed 2026-06-28.

Baptiste Vasey, Myura Nagendran, Bruce Campbell, David A. Clifton, Gary S. Collins, Alastair K. Denniston, Livia Faes, Bart Geerts, Mudathir Ibrahim, Xiaoxuan Liu, Bilal A. Mateen, Peter McCulloch, and DECIDE-AI Expert Group. Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: Decide-ai. *Nature Medicine*, 28: 924–933, 2022. doi: 10.1038/s41591-022-01772-9.

World Health Organization. Ethics and governance of artificial intelligence for health: Who guidance, 2021. URL https://www.who.int/publications/i/item/9789240029200. Accessed 2026-06-28.