MEDOS INTELLIGENCE LAYER

AI as a system, not a feature.

A recommender-first clinical & operational intelligence layer — from bedside vision safety to a C-suite decision twin, on one governance spine.

Shipped surfaces are sandbox-verified. Roadmap is labeled honestly — never overclaimed.

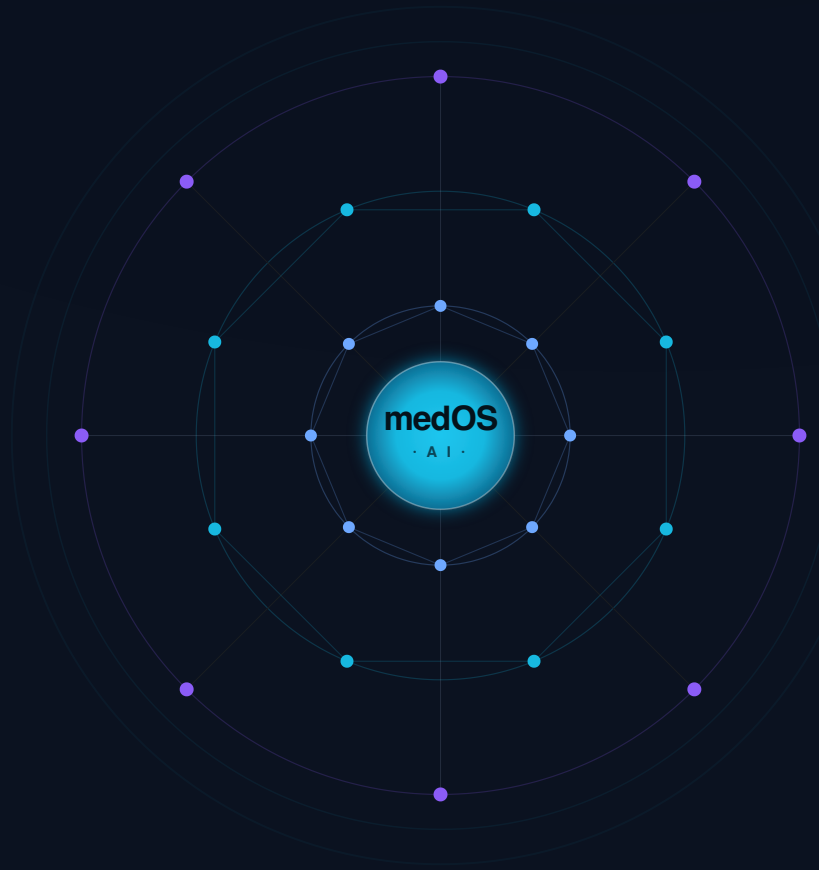

THE PROBLEM

AI is sold as a checkbox. Hospitals get a junk drawer of demos.

🧩 Every tool is its own silo

Separate login, separate data assumptions, separate audit trail. Nothing shares a spine.

✍️ Most "AI features" quietly write to the chart

No human gate, no kill switch, no provenance — an autonomous write you only discover later.

📤 PHI leaks into vendor cloud prompts

And the safety logic is hardcoded — not tunable per hospital, per ward, per scheme.

🔌 Nothing connects the three altitudes

Bedside safety, the ops backlog, the boardroom decision — three unrelated products.

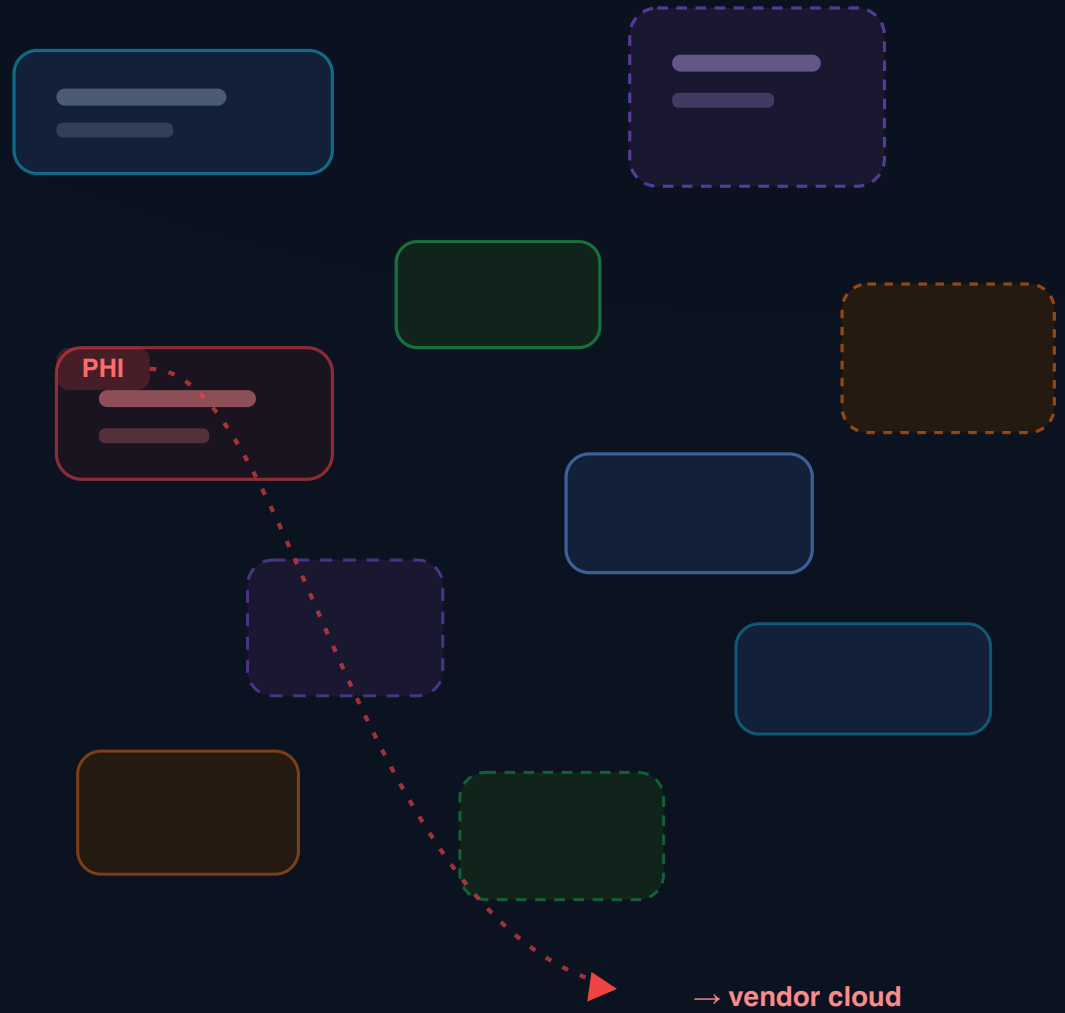

THE THESIS

One spine, five pillars: AI that proposes, humans that dispose.

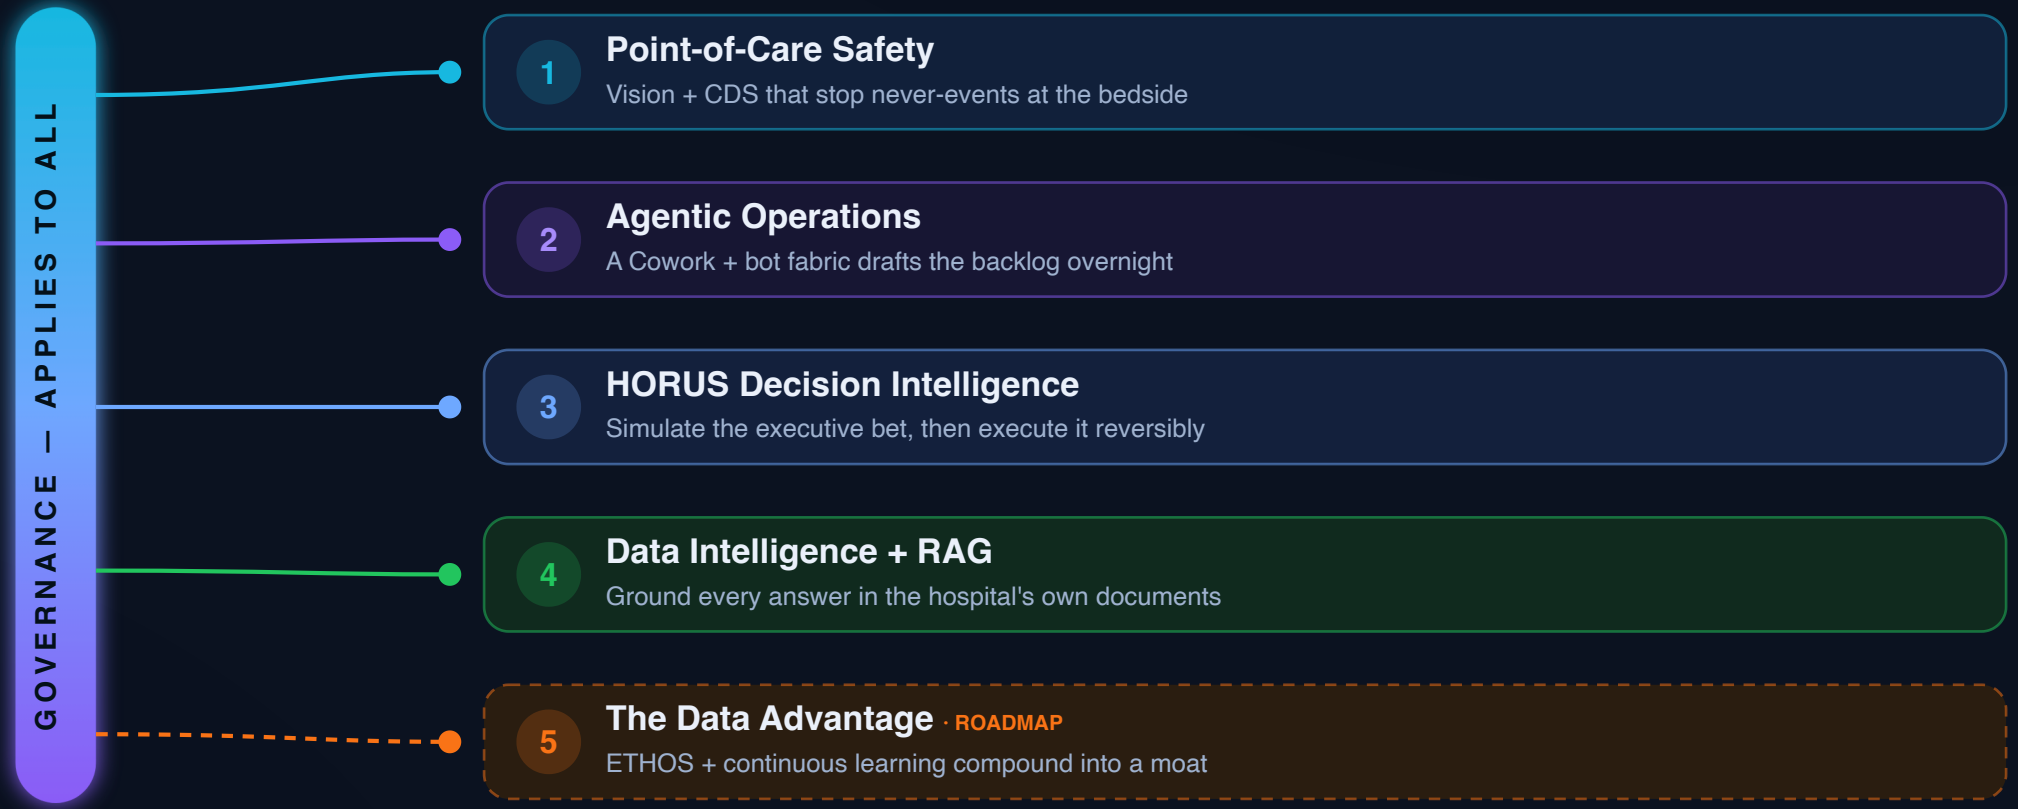

◆ MEDOS · AI

PILLAR 1 · POINT-OF-CARE SAFETY

Vision + CDS that stop never-events before they happen.

Surgical Instrument Count AI SHIPPED · MAIN

YOLO vision over 22 AORN instrument types across 4 count phases — so no sponge or clamp is left behind.

Configurable CDS Vital-Signs Engine SHIPPED

NEWS2 / MEWS / qSOFA-class rules fire on *every* `recordObservation()` write — screening, nursing, anesthesia, pharmacy, HL7v2, FHIR.

Medication Safety Verdict Agent READY · BRANCH

Not yet on main: the one vision module that already has its orchestrator handler + e-MAR critical-block wired (OCR+LLM → safe / caution / warning / critical).

Bedside verification suite READY · BRANCH

Wristband ID · med-tray match · blood-bag ABO/Rh · specimen label · vitals-monitor OCR — each gates the next step via tunable `policy_gates`.

Brilliant AI Screening PARTIAL · DEFAULT-OFF

Chief complaint + vitals + deterministic red-flag check → nurse-confirmed ESI 1–5 triage.

◆ MEDOS · AI

PILLAR 2 · AGENTIC OPERATIONS

A Cowork fabric that drafts the hospital's backlog overnight.

Agentic substrate SHIPPED · DEPLOYED

`cowork_agents` identity · `cowork_proposals` envelope · `automation_bots` runtime · AcknowledgementInbox — ~70–80% of an agent product, deployed.

Recommender Delegation Bridge SHIPPED

Every output is a draft. Accept / Edit / Reject is the actuation seam *and* the training label — never an autonomous chart write.

Backend Agent Runner SHIPPED

Server-side LLM + tools fire on cron / event sweeps — work appears with no browser tab open.

Patient Automation Hub SHIPPED

One per-patient timeline unifying 6 automation / task / queue sources: "which bots are operating now, and what's next."

Roadmap agents DESIGNED

Collections (Promise-to-Pay) · Disruption-Recovery · Supply-Chain charge-leakage · Ambient Scribe · OPD Autopilot — all recommender-gated.

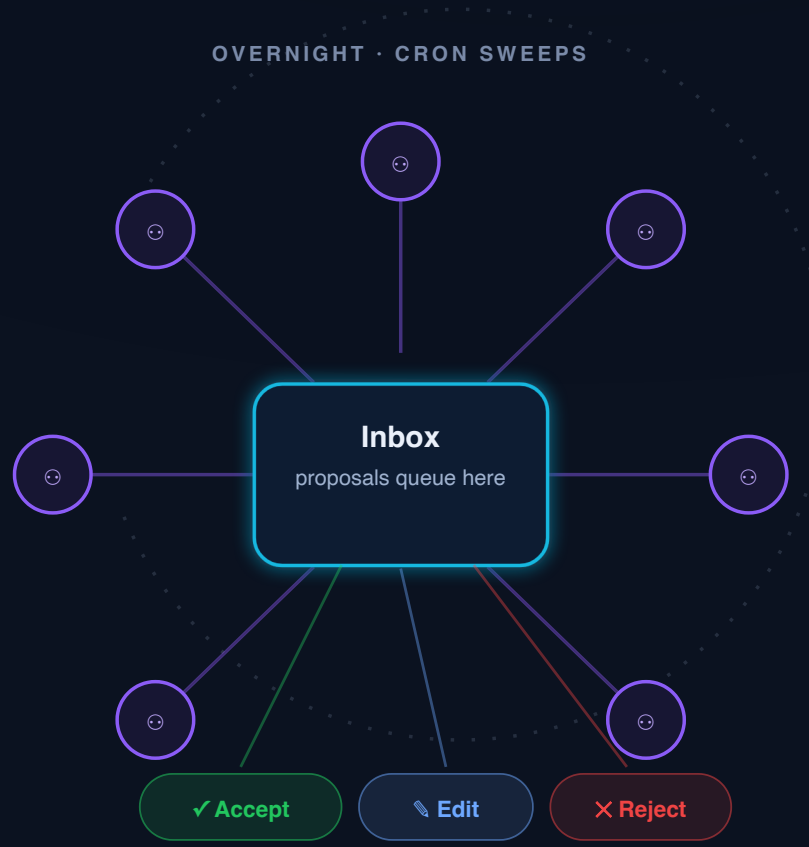

…the gesture is the training label

◆ MEDOS · AI

PILLAR 3 · HORUS DECISION INTELLIGENCE

Simulate the executive bet. Then execute it — reversibly.

HORUS Decision Twin PARTIAL · P0–P5

Discrete-event + Monte-Carlo engine returns P10 / P50 / P90 risk bands for C-suite bets (step-down unit, cath lab, payor scheme). **On seed data today**; live per-metric calibration remains.

Seven Domain Agents PARTIAL

Demand · Capacity · Workforce · Supply · Reimbursement · Quality · Compliance — designed to calibrate from the hospital's own history. v1 on seed data; 3 of 7 wired to the live-data boundary.

Initiative Optimizer + Decision Guide PARTIAL · IN-LENS

Runs client-side on seed data: readiness score, investment memo, phased roadmap — worked example recommends **4 → 7 → 10 beds** over 6 months.

HORUS Map Lenses SHIPPED · MAIN

7 live operational views: 3D Atlas · SEIR surveillance · capacity · OPD queue · FP&A margin · RUDS threat map.

Initiative Executor SHIPPED · DEFAULT-OFF

A signed decision fans out to tasks, acks, gate-flips, cron, Jira/GitHub via 8 action_kinds — closing the gap where approvals dead-ended at an audit row.

SIMULATE — P10 / P50 / P90

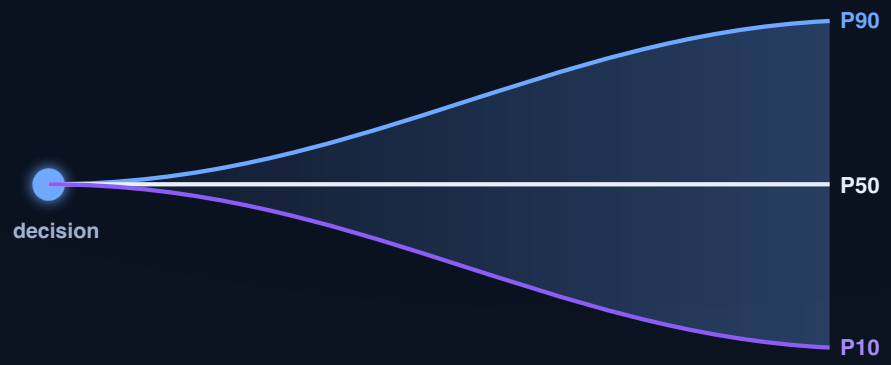

SIGN-OFF — THEN EXECUTE

tasks · acks · gate-flips · cron · Jira / GitHub

every action carries a rollback snapshot

◆ MEDOS · AI

PILLAR 4 · DATA INTELLIGENCE + RAG

Every recommendation grounded in the hospital's own documents.

Document Intelligence + RAG PARTIAL · P0

Activate Excel / Confluence / folders into a private, de-identified corpus; ask a local self-hosted model that *cites its sources*.

Smart Intake Inbox PARTIAL

Drop a folder; AI identifies the patient or hospital system and proposes where to file. AI-guessed HN capped at 0.85 — it can never auto-file.

Data Activation Center (HDAP) SHIPPED

`/admin/hdap`: one hub for connectors, file activation, AI analysis + vertical presets (Teaching Hospital / Wellness Center).

AI Form Builder & Form-Fill PARTIAL

Fills or builds a form bound to 4,535 grounded operations (2,421 actions + 2,114 entities) across 13 services — no fabricated endpoints.

5-Domain Recommendation Engine PARTIAL

Ranked diagnosis / order-set / code / quick-pick on a "your history > department > hospital" signal. Model loaded on PH; engine still disabled (flag unset + MVs not applied).

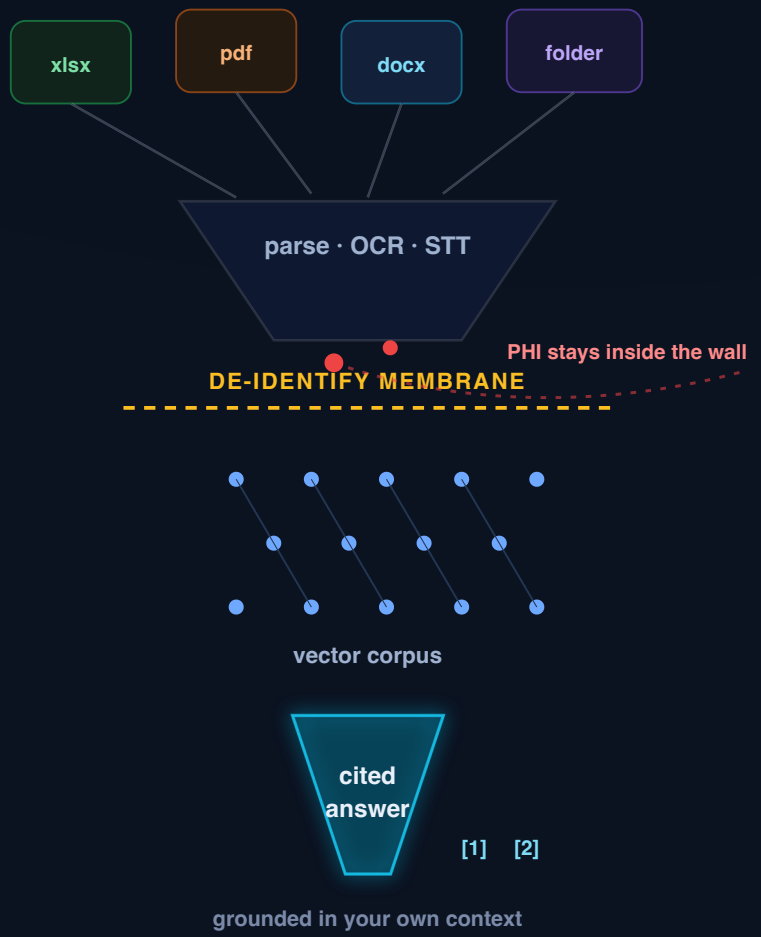

ROADMAP

PILLAR 5 · THE DATA ADVANTAGE — ROADMAP

Trained in-region on your own outcomes — the part competitors can't copy.

Hospital Data Advantage DESIGNED

Calibrate the twin on 10M+ of the hospital's own outcome histories; ground every answer in its own M&M notes, committee minutes, SOPs.

Central Trajectory Engine — ETHOS DESIGNED

A decoder-GPT over tokenized Patient Health Timelines, *retrained in-region* — published weights are a non-redistributable MIMIC derivative.

Reported research figures ETHOS ON MIMIC

Inpatient-mortality AUC **0.921** · hospital 30-day readmission AUC **0.749** — research figures, not medOS-measured. (0.807 is ICU readmission.)

ETHOS-NG DESIGNED

RoPE + FlashAttention-2 (2,048 → 8,192 ctx) · continuous value embeddings · survival head · conformal calibration · QLoRA / GGUF edge export.

Continuous-Learning (intelligence-trainer) DESIGNED

Retrain → eval on a frozen golden set → gate → promote. Corpus opt-in (default FALSE), K=20 k-anon, in-region.

PATIENT HEALTH TIMELINE → FUTURES

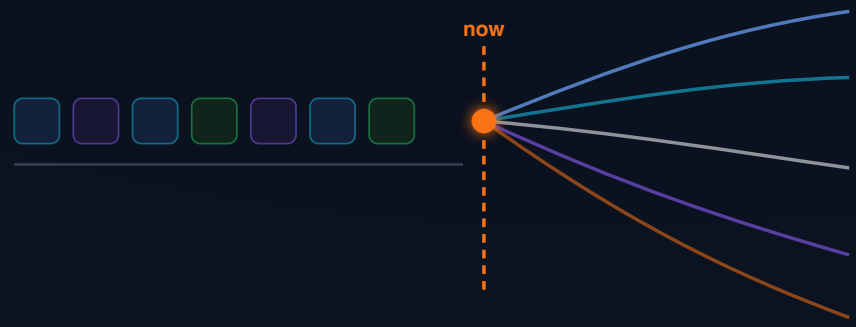

CONTINUOUS LEARNING

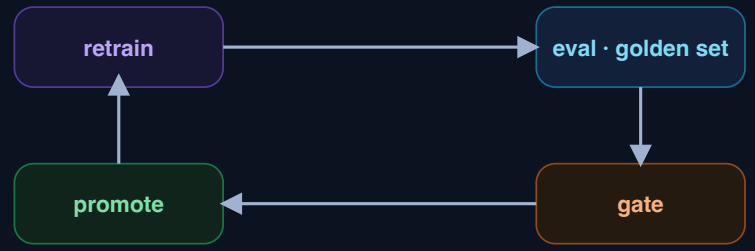

promotion to a decision surface is human-gated

◆ MEDOS · AI

PROOF · WHAT'S REAL TODAY

Shipped, sandbox-verified, and running on the PH demo.

Surgical Instrument Count AI

22 AORN instrument types · 4 count phases. On `main`.

Configurable CDS Engine

Fires on every `recordObservation()` write. `/admin/cds-rules`.

HORUS Map Lenses

7 live operational views shipping on `origin/main`.

Initiative Executor + Rollback

8 action_kinds · append-only ledger · one-click undo.

Cowork Substrate + Bridge + Runner

~70–80% of an agent product, deployed.

Patient Automation Hub

6 automation / task / queue sources unified per patient.

In-Platform CRM Spine

Contacts · deals · campaigns · tickets · consents — applied + live.

Inventory Ledger + Rules Engine

3,892 balances live · consume → charge, no leakage.

Vision Backend + Auto-Repair T1

Selector live on `main` for Surgical Count; same adapter ready for branch modules. Self-healing restart.

Every card here is **shipped on main** or **applied + verified live** — designed items are deliberately kept off this grid.

◆ MEDOS · AI

PROOF · THE HONEST STATUS LEDGER

We label shipped, partial, and designed — out loud.

● SHIPPED

- Surgical Count · CDS Engine
- HORUS Map Lenses (7)
- Initiative Executor + Rollback
- Cowork substrate · bridge · runner
- Patient Automation Hub
- In-Platform CRM Spine
- Data Activation Center (HDAP)
- Inventory ledger + rules
- Vision backend selector · Auto-Repair T1
- Agent audit / provenance

◐ PARTIAL — BUILT, GATED OR UNAPPLIED

- HORUS Twin (P0–P5, seed data)
- 7 Domain Agents (3 live) · Optimizer
- Boundary contracts (9 metrics)
- Brilliant Screening · OPD Autopilot
- Med-Rec v2 · bedside vision suite (branch)
- Document Intelligence / RAG · Intake Inbox
- AI Form Builder · 5-Domain Reco Engine
- RUDS (P1–3) · Voice Attestation
- PM-Sync (2 of 8) · AHP gate

○ DESIGNED — ROADMAP

- Hospital Data Advantage
- ETHOS · ETHOS-NG
- intelligence-trainer (continuous learning)
- Collections · Disruption-Recovery agents
- Ambient Clinical Scribe
- Unified Clinical Assistant
- AI Problem List · Medication AI Planner
- Nav Next-Best-Action · AI Training Corpus
- Multilingual AI · Auto-Repair T2

Shipped Partial Designed

The gap between **partial** and **shipped** is usually a migration to apply or a flag to flip — **not net-new code**.

◆ MEDOS · AI

GOVERNANCE · THE SPINE

Six invariants keep every output on the safe side of the device line.

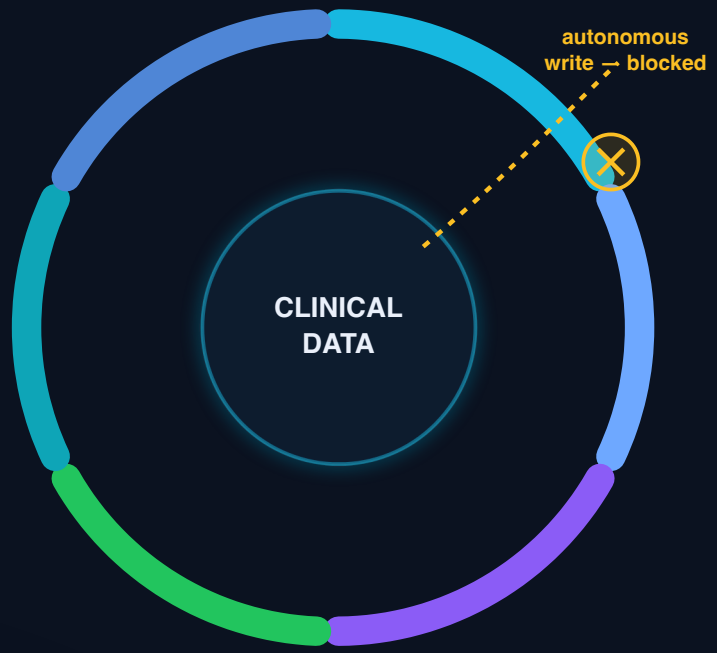

1 · Deterministic safety first
A red-flag / escalate verdict or CDS BLOCK always wins — AI output is suppressed, never merged or downgraded.

2 · No autonomous clinical or payment writes
Ever — not even behind a flag. "Auto" is statically rejected at config-write time.

3 · Accept / Edit / Reject is the only actuation seam
And it doubles as the training label. Agents propose into the Inbox, never a system-of-record table.

4 · Human sign-off as the SaMD firewall
No twin recommendation becomes a "decision" without four named roles signing off.

5 · PHI-safe by design
Modeled at the DB-grant level (3-plane), per-plane grants + RLS hardening in progress; FORBIDDEN_KEYS reject PHI at ingest; in-region inference, no PHI in prompts.

6 · Off-by-default + kill switch + reversible
Disabled means zero behavior change; every executor action carries a before-snapshot for one-click rollback.

◆ MEDOS · AI

THE MOAT

Swappable over any HIS. Trained on data only you have.

Five boundary contracts REGISTERED · 9 METRICS

A metric-RPC registry, foreign-HIS sync adapter, ERP/HRM facts feed, recommender proposal seam & market-evidence feed — designed so the decision layer runs over medOS today or a foreign HIS (HOSxP / incumbent EMR) tomorrow by swapping the reader. Foreign-HIS swap not yet demonstrated.

Codes are the interlingua

The AI reasons over universal ICD / ATC-TMT / LOINC tokens — so a new language translates only the narration shell, not the reasoning layer.

The data advantage compounds

Every answer grounds in the hospital's own RAG corpus; the trajectory model retrains in-region on the customer's own consented outcomes.

Retrain — don't redistribute

Each deployment's model is population-calibrated and ours to distill to the edge — a moat a generic API vendor structurally cannot copy.

One spine + one substrate

New capabilities mount as rows and seeds, not forks.

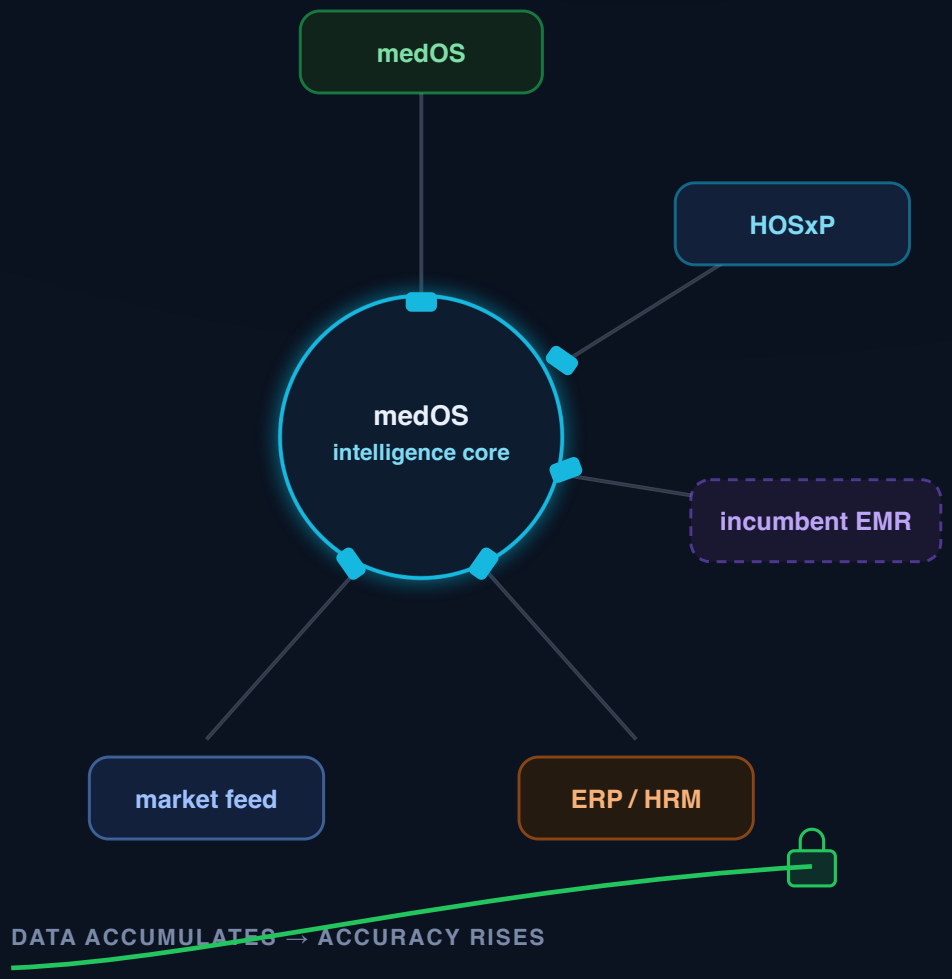

THE CLOSE

Buy the system, not the demo.

TODAY

Bedside never-event safety · a deployed agent fabric · a working twin with 7 live map lenses · a grounded RAG layer — all sandbox-verified.

TOMORROW

ETHOS trajectory generation + continuous learning compound your own data into accuracy no API vendor can match.

ALWAYS

Deterministic-safety-first · no autonomous writes · human sign-off · PHI behind the DB wall · off-by-default with a kill switch.

Start with one capability — CDS or Surgical Count →