

A HOME NURSE ROBOT

The nurse that doesn't leave at the end of a shift.

A mobile in-home care device that dispenses medication it can prove it delivered, monitors the vitals that flag a crisis, and summons help when no one else is there — engineered so the parts that must never be wrong are the parts that cannot drift.

VERIFIED DISPENSING

CONTINUOUS VITALS

INDEPENDENT ESCALATION

One device, three machines, and the single decision the whole design rests on.

A home nurse robot is not one machine. It is the union of three with radically different tolerances for error, sharing one body: a general-purpose service robot that fetches and navigates a cluttered home; a vital-signs monitor that must know a real anomaly from a bad reading; and a medication dispenser and emergency-escalation device where the only acceptable behaviour is correct behaviour.

The founder's constraint — *“it can't drift, it has to do everything correctly”* — is not answered by making the AI better. It is answered by removing the AI from the paths where correctness is non-negotiable. A learned model plans and manipulates the untidy world; a separate, deterministic, independently-verifiable controller owns every actuator motion that could harm a person, and owns the two life-critical decisions: **what pill is released** and **whether to escalate to 911**. The intelligent system decides *what to attempt*; the deterministic systems decide *whether it is allowed* — and they are the only ones wired to the muscles.

This is standard practice in every field where software can kill: avionics, industrial safety, infusion pumps. It is why the rest of this document holds together for the two readers it is written for at equal weight — the investor pricing a market, and the engineer checking the part numbers.

“Not a promise that the AI is perfect — a system in which the AI is not permitted to be the single point of failure for anything that can hurt someone.”

CONTENTS

I THE INVESTMENT CASE

The problem, and why now	04
Market, built two ways	07
Competitive landscape	09
Business model & reimbursement	11
The raise	13

II THE MACHINE

System architecture	16
The brain: two, kept apart	18
Actuation & the dispenser	20
Sensing & the honest accuracy	22
The safety case	24
Bill of materials	26
The virtual twin	27

III THE MOAT

FDA classification	29
Pre-Sub, SaMD & the PCCP	31
Privacy & liability	33
Phased path & the numbers	35

APPENDIX

Sources & verification ledger	37
-------------------------------	----

ON THE NUMBERS

Every load-bearing figure in this dossier is verified against a named source read for this document, or written as an explicit range with its basis stated. Where reputable sources disagree, both figures are shown and the conservative one anchors the case. A market size is a claim with a source or it is a guess — this document labels which.

A five-figure care bill, a shrinking workforce, and a machine positioned exactly where no incumbent stands.

The United States is aging on a curve set by births that have already happened, while the workforce that would provide the care is moving the opposite direction. The gap is absorbed today by unpaid family labour valued at **\$600 billion a year**. That is the labour this device is built to offload — and it is measured against a private nursing-home room at **\$127,750 a year**, not against zero.

No fielded product combines mobile presence, *verified* medication dispensing, vitals monitoring, and independent emergency escalation in a single in-home device. Dispensers don't move or monitor; companions remind but can't act; monitors escalate but can't dispense. The white space is real — and it is empty because it is hard and regulated, which is precisely why clearing it is the moat.

58 → 82M US ADULTS 65+, 2022→2050 (+42%)	9.7M DIRECT-CARE OPENINGS TO FILL, 2024–2034	\$600B UNPAID FAMILY CAREGIVING, PER YEAR
\$25–30B BOTTOM-UP SERVICEABLE MARKET (SAM)	Class II FDA PATH — 510(K)/DE NOVO, NOT PMA	\$3.5M PRE-SEED RAISE, MILESTONE-GATED

WHAT IT IS

A mobile home device on a deliberately slow wheeled base, with a light manipulation arm for fetch-and-carry, a locking unit-dose medication dispenser, contactless and contact vitals, and a caregiver telepresence window — the “camera-in-eyes.” The high-level autonomy runs a π0-class vision-language-action model on a Jetson Thor-class module; the life-critical decisions run on separate, deterministic microcontrollers that the AI can only *request*, never command.

WHY IT IS FUNDABLE NOW

Three curves crossed in the last 24–36 months: physical-AI foundation models matured, edge compute got small and powerful enough to run the safety stack locally, and remote care normalised and got *reimbursed*. The raise is sized to reach a working prototype and an FDA Pre-Submission — the two events that convert regulatory risk from unpriced to priced and make a Series A rational.



I

PART ONE — THE INVESTMENT CASE

The demand curve is not a forecast. It is birthdays that have already happened.

Market, timing, competition, the business model, and the raise — every load-bearing figure sourced or ranged, the arithmetic shown.

THE INVESTMENT CASE

Two curves are crossing, and one of them cannot be argued with.

The population age 65 and older is projected to grow from **58 million in 2022 to 82 million by 2050** — a 42% increase — its share of the total population rising from 17% to 23%. By 2034, for the first time in US history, older adults will outnumber children. The sharper number is the 85+ “oldest-old,” the cohort that actually needs help with medication and monitoring: it more than **doubles, from 6.5 million to 13.7 million by 2040**.

The people who would provide the care are not there. The direct-care workforce is the fastest-growing occupational group in the country and it cannot be filled: **roughly 9.7 million openings between 2024 and 2034**, most from turnover, not new demand. The ratio of people 85+ to those in the peak caregiving years moves from about **1-to-7 today toward 1-to-3 by 2050**. There are not enough hands.

The gap is absorbed by unpaid family, and its scale is the clearest proxy for unmet need. AARP values that work at **\$600 billion a year** — 36 billion hours in 2021, more than all US out-of-pocket healthcare spending that year. That \$600 billion is the labour this device is designed to offload.

THE COST GAP MAKES AGING-IN-PLACE THE ONLY SUSTAINABLE ANSWER

A device that substitutes for even part of the care labour is measured against a five-figure annual bill, not against zero. This is the price umbrella under which a hardware-plus-subscription product profitably sits.

CARE SETTING, US 2024	ANNUAL MEDIAN
Nursing home, private room	\$127,750
Nursing home, semi-private	\$111,325
Home health aide	\$77,792
Homemaker services	\$75,504
Assisted living community	\$70,800

Source: Genworth / CareScout 2024 Cost of Care Survey (15,000+ provider surveys).

THE WEDGE: ADHERENCE

Medication non-adherence costs the US health system an estimated **\$100B–\$290B a year**; the broader “non-optimized medication therapy” figure is **\$528.4B (2016)**, associated with ~275,000 deaths annually. This dossier keeps the two distinct.

More than **100,000 Americans** are estimated to die each year from non-adherence; up to **10% of hospital admissions** are attributed to it, and it roughly *doubles* readmission within a month of discharge — the exact event that drives avoidable cost.

WHY NOW, AND NOT 2016

- 1 — Physical-AI matured.** Vision-language-action models moved general manipulation from hand-coded routines to models that generalise across tasks.
- 2 — Edge compute got good enough.** Running the perception-and-safety stack locally — responsive and privacy-preserving — did not exist at this price and power in 2016.
- 3 — Remote care got reimbursed.** COVID permanently shifted expectations, and CMS built a durable reimbursement scaffold for remote monitoring. A camera in the home and a vitals stream are no longer strange; they are billable.

A disciplined SOM of \$150–400M, not one percent of a trillion-dollar mirage.

A market slide is the first thing a serious investor fact-checks, so this section shows its work. The right unit is a household with an older adult who lives at home, manages a real medication regimen, and can pay — directly, through a family member, or through a payer. Every step is sourced or labelled an assumption.

THE BOTTOM-UP FUNNEL

LAYER	HOUSEHOLDS	ANNUAL VALUE*
TAM	35–40M	\$60–70B
All US 65+ households, at home, with a care/medication need		
SAM	14–16M	\$25–30B
75+ at home ×90% ×80% regimen = ~18–19M individuals; payable households		
SOM (5-yr)	75–225k units	\$150–400M
~0.5–1.5% of SAM households, ARR at scale		

*Blended ~\$1,800/household/yr — hardware amortised into subscription plus the RPM reimbursement layer. Modelled below Hero's bare-dispenser stack plus a management layer, and well below the human care it displaces. Sources: Census P25-1144, PRB, ACL; assumptions disclosed in §2.

THE HONEST CAVEAT

No home robot has yet proven mass-market adoption at consumer price points; the US smart-dispenser market is only ~\$0.8B today. That is evidence the *category is nascent*, not that demand is absent. The SOM is an execution-and-adoption bet. Demographics guarantee the demand exists; they do not guarantee this form factor captures it.

TOP-DOWN ANCHORS (CONTEXT, NOT TAM)

US home healthcare	\$110→168B
2025→2030, 8.9% CAGR — Grand View / M&M	
Remote patient monitoring	\$27.7→56.9B
2024→2030, 12.7% (conservative anchor) — M&M	
Smart pill dispenser	\$2.1→4.2B
2024→2030; US share ~\$0.8B — Strategic Market Research	
Unpaid family caregiving	\$600B
The labour substituted — AARP 2023	

ON RPM'S SPREAD

Three RPM sources disagree by more than 2× on today's size and ~7× on the out-year. This dossier uses the \$27.7B→\$56.9B anchor and *states the range* rather than blending. An investor who checks will find all three — honesty about the spread is more credible than false precision.

The cross-check: the bottom-up SAM sits between “all of RPM” and a small slice of home healthcare — and an order of magnitude above the entire current dispenser market. That is the sanity check an investor wants.

A technical reader knows these companies. Strawmen are worse than useless.

Every entry is built on verified funding and product facts, and is honest about where each competitor is genuinely strong. Each incumbent solves one slice — dispensing, or the panic button, or companionship, or monitoring — and none integrates them into a single mobile agent.

ASSISTIVE HOME ROBOTS (MOBILE)

Labrador Systems — the Retriever. The closest true peer: a wheeled self-driving robot that carries things for people with mobility challenges. It retrieves and transports; it does *not* dispense medication, take clinical vitals, or escalate.

\$3.1M seed, 2022 co-led by Amazon Alexa Fund & iRobot Ventures — informative about how hard consumer home robotics is to fund.

Diligent Robotics — Moxi. A genuinely successful service robot, but adjacent: fetch-and-deliver logistics *inside hospitals*, non-patient-facing. ~\$50M+ total Proves service robots get deployed and paid for in healthcare — but Diligent is not building for the home.

COMPANION / SOCIAL ROBOTS

Intuition Robotics — ElliQ. The companionship category leader: reduces loneliness, runs check-ins, gives medication reminders — but no physical dispensing, no clinical vitals, no mobility, no escalation. ~\$83M total Its 3,000+ units and first Medicare-MCO agreement *prove the payer/agency channel* this product would also pursue.

Catalia Health — Mabu. A tabletop conversational health companion focused on adherence *conversations*, not dispensing. A cautionary point: a talking head that reminds but cannot verify or act has struggled to define a durable business.

MEDICATION DISPENSERS (THE REAL NEAR-TERM COMPETITION)

Hero Health. The market-defining smart countertop dispenser: stores, sorts, dispenses, tracks, reorders, alerts. \$12M Series A Pricing ~\$360/yr up to \$59.99/mo + ~\$99 activation; HIPAA-compliant, FDA-registered. The feature floor a robot must clear — *but a stationary box: it cannot move to the person, take vitals, or respond to a fall.*


MedMinder, Livi, MedaCube, PERS incumbents.

MedMinder ~\$50M and the dispenser tier occupy the locked-per-dose category; Philips Lifeline, Medical Guardian and Life Alert own the “call for help” button. Each solves one slice.

MONITORING & BEHAVIOURAL-AI ELDERTECH

Sensi.AI (audio care copilot, ~\$31M), **CarePredict** (wearable behavioural AI, ~\$28M), **Aloe Care Health** (voice PERS + family network, ~\$30M). Monitoring-and-alerting platforms — not physical devices that dispense or act.

No one credibly combines mobile presence, verified dispensing, vitals, and escalation in one in-home device.



Two axes: *can it physically act in the home* (dispense, move to the person) × *does it monitor and escalate clinically* (vitals, anomaly detection, emergency call). The upper-right quadrant is empty.

The white space is not empty because no one thought of it. It is empty because it is hard and regulated.

A device that dispenses medication *and* monitors vitals *and* autonomously calls for help is, almost certainly, an FDA-regulated medical device — on a clearance pathway that companion robots and even dispensers have partly sidestepped by staying “reminders, not treatment.”

The moat *is* the difficulty and the regulatory clearance. Which is exactly why the raise funds a working prototype and an FDA pre-submission — the two milestones that de-risk the Series A. The market strategy and the regulatory strategy are one story: the white space is defensible precisely because clearing it is hard.

THE DIFFERENTIATED CLAIM

A product that verifies a dose was dispensed and taken, checks the vitals that would flag a problem, and escalates when something is wrong — attacking the readmission-and-non-adherence cost driver in a way no single incumbent does.

Robot-as-a-service: hardware amortised into a recurring, partly-reimbursed subscription.

The device is capital-intensive to build, so a pure hardware sale would put a large sticker price in front of a price-sensitive buyer. The category has already answered this: Hero, ElliQ, and the dispenser incumbents all sell subscription, not hardware. The recommended model follows — low or zero upfront cost, a monthly fee bundling the device, software, monitoring, caregiver access, support and warranty.

THE BLENDED-MARGIN TARGET

At volume, in the range typical of connected-medical-device-plus-service businesses — roughly **50–65% blended**, with the subscription and RPM layers well above that and the hardware below it. The exact figure is a function of the delivered BOM and is stated in the financial model, not asserted here.

THREE PAYERS, IN ORDER OF NEAR-TERM REALISM

01

Families, out of pocket

The beachhead. The adult child managing a parent's care — high-intent, measuring cost against a five-figure facility bill or their own unpaid labour. A **\$100–200/month** subscription is credible against the alternatives. Where ElliQ and Hero found first traction.

02

Medicare Advantage & MCOs

The scale channel. Capitated plans profit when a member *avoids* a hospitalisation or facility stay. Not speculative: Intuition Robotics already signed its first MCO agreement, and ElliQ is distributed free through state Offices for the Aging. The channel is proven; the product earns its place with outcomes data.

03

LTC insurers & employers

LTC insurers face the same \$128K/yr nursing-home curve and have a direct incentive to fund aging-in-place. Employer caregiver benefits (Homethrive, Cariloop, Wellthy) are a fast-growing B2B2C route.

UNIT ECONOMICS — THE THREE LEVERS

COGS falls with volume. The target is a device COGS that a \$100–200/mo subscription recovers within ~18–24 months, after which the unit is gross-margin-positive across a multi-year service life.

Subscription margin compounds. Once hardware is amortised, the marginal cost of another month of monitoring is small — software-like margins over a 3–5-year life in the home.

RPM reimbursement stacks on top. Incremental revenue against a largely fixed device cost — the closest thing to free money in the model, and the reason payer distribution is the prize.

THE HONEST CONSTRAINT

RPM billing requires an established patient relationship, a billing clinician, the 16-day data-capture rule, and clean HIPAA handling of the vitals and camera stream. It is also a scrutinised billing area — CMS has flagged fraud in the category.

This is revenue that must be *earned through partnership and clean operations*, not switched on. But it is real, recurring, government-backed revenue — the single most important reason this is a healthcare business and not a consumer-electronics gamble.

CMS already pays for remote physiologic monitoring. A device that continuously takes vitals is exactly the hardware these codes were written for.

CPT CODE	WHAT IT COVERS	~2025 AVG.
99453	Initial setup & patient education (one-time)	\$19.73
99454	Device supply + data transmission / 30 days (≥16 days)	\$43.02
99457	First 20 min/mo clinical management + interaction	\$47.00
99458	Each additional 20 min/mo (add-on)	\$38.64
99091	Physician collection & interpretation, 30 min	\$52.00

Sources: CMS CY2025 PFS Final Rule (conversion factor \$32.35), via Smart Meter, HRS, Tenovi, PYA. Non-facility amounts; vary by geography. Rates stepped down modestly from 2024 as the PFS factor fell ~2.8%.

\$100–160

PER MONITORED PATIENT, PER MONTH, RECURRING, PAID BY MEDICARE

A single monitored patient can support ~\$100–160/month in RPM billing (99454 + 99457, often + 99458), provided the device captures ≥16 days of physiologic data per month and a qualified clinician does the management.

That reimbursement can subsidise the subscription — the effective out-of-pocket cost to the family can sit substantially below the true service value while the business collects on both sides. It anchors the product firmly in the clinical, reimbursed world rather than the discretionary-gadget world, which is where durable healthcare businesses live.

\$3.5M pre-seed. 18–24 months. Two hard milestones.

The working recommendation was a \$2.5–4M pre-seed. After pressure-testing against real rounds in home robotics, eldertech and medtech hardware, the verdict: **the structure is correct, and the number is defensible** — target the upper half, \$3.5–4M, framed as a milestone-gated pre-seed rather than an open-ended one.

WHAT COMPARABLE COMPANIES ACTUALLY RAISED (ALL VERIFIED)

COMPANY	STAGE / ROUND	AMOUNT	WHAT IT BOUGHT
Labrador Systems	Seed, 2022	\$3.1M	Consumer home assistive robot — <i>closest hardware peer</i>
Diligent Robotics	Seed → total	\$3M → ~\$50M	Hospital service robot (Moxi)
Intuition Robotics	early → total	\$6M → ~\$83M	Companion robot (ElliQ)
Hero Health	Series A, 2018	\$12M	Smart med dispenser
MedMinder	growth	~\$50M	Dispenser + adherence
Sensi.AI / CarePredict / Aloe	growth	\$28–31M ea.	Eldertech monitoring platforms

1 — UNDER-CAPITALISED AT SEED

Labrador — the single closest hardware analog — raised only **\$3.1M at seed**, even with Amazon and iRobot behind it. Capital-efficient hardware-robotics seed rounds in this exact category cluster at \$3–4M.

2 — THE BIG ROUNDS CAME LATER

The \$30–83M rounds happened at Series B/C, *after* a shipping product and traction. That capital is the next milestone, not this one — the prize the pre-seed is designed to unlock.

3 — WHAT DE-RISKS THE A

(a) a working prototype that demonstrates the integrated capability, and (b) an FDA pre-submission that turns “is this even clearable?” into a documented plan. Every pre-seed dollar points at those two events.

WHY THE UPPER HALF OF THE RANGE

This product carries *two* capital-intensive burdens at once — building physical hardware *and* funding regulatory groundwork — where Labrador carried only the first. The extra ~\$1M buffer buys the regulatory pre-submission work and a margin against hardware iteration. Going in at the bottom risks running out of runway before the FDA pre-sub is complete, forcing a bridge on weak terms.


ILLUSTRATIVE USE OF THE \$3.5M

Engineering & working prototype 45–55%




Core team, prototype BOM, the VLA/autonomy and safety stack

Founding team & operations 20–25%




Founder salaries, key early hires, G&A over 18–24 months

Regulatory & clinical groundwork 15–20%



FDA Pre-Sub prep, regulatory/quality lead, early hazard analysis & safety case

Pilot & evidence 5–10%



A small monitoring/telepresence pilot — the first outcomes signal for the payer story

THE SERIES A THESIS THIS SETS UP

“We have a working integrated prototype, a documented FDA pathway from our completed pre-submission, a monitoring/telepresence product generating early RPM revenue and payer interest, and the demographic and cost tailwinds are only strengthening. Fund the clearance and the scale.”

That is a raisable Series A — and every dollar of the pre-seed is spent to be able to say it. In medical devices, revenue and clearances compound: you buy the hard clearance with the money and credibility the easy ones generate.

THE MILESTONE LADDER — GO-TO-MARKET AND REGULATORY PHASING ARE ONE SEQUENCE

PHASE 1 • BEACHHEAD

Monitor + telepresence + reminders

Ship the RPM-eligible, caregiver-valuable parts without waiting on full dispensing clearance. Puts robots in homes, generates first reimbursement and outcomes data, validates the payer channel. The lowest-regulatory-burden entry — and the cleanest 510(k) predicates.

PHASE 2 • ESCALATION

Validated anomaly detection + tiered escalation

With Phase-1 data in hand, add the monitored-escalation function under the human-in-the-loop model, cleared via the De Novo pre-aligned with FDA in the Pre-Sub. This is where the category-defining regulatory asset gets created.

PHASE 3 • DISPENSING

Autonomous Rx dispensing

The hardest, most valuable function — the full Remote Medication Management System — clears last, funded by Phase 1/2 revenue and de-risked by the FDA relationship, the QMS, and the safety track record built along the way.

Why this is fundable, not a retreat: each phase is a real product with real revenue and a real clearance; each de-risks the next; and the sequence spends the cheap, fast clearances first to earn the right to attempt the expensive, slow one. The pre-seed is sized to reach a shippable Phase-1 prototype *plus* an FDA Pre-Sub — not to reach a clearance, and the pitch does not claim it does.

II

PART TWO — THE MACHINE

**Make the parts that can
tolerate learning smart.
Make the parts that
cannot dumb,
deterministic, and
provable.**

The architecture, the brain, the actuation and the sealed dispenser, the sensing and its honest accuracy limits, the safety case, the bill of materials, and the virtual twin. Every load-bearing spec verified against the vendor's or standards body's own page.

3

INDEPENDENT
PROCESSORS

0


LIFE-CRITICAL ACTS IN THE
"LEARNED" COLUMN

1

DOSE PER AUTHORISED CYCLE,
MECHANICALLY

The intelligent system decides *what* to attempt. The deterministic systems decide *whether it is allowed* — and they are the only ones wired to the muscles.

Three machines, one body. Collapsing them into a single learned policy is the mistake that makes the “can’t drift” objection fatal. Kept separate, the parts that can tolerate learning get to be smart, and the parts that cannot are dumb, deterministic, and provable.



The two-brain split, drawn to the interface. The dispenser is deliberately a *third* processor, not a feature of the supervisor, because its correctness argument is independent: it will not release a dose the schedule does not authorise even if both the cortex and the supervisor tell it to.

The cortex never drives an actuator directly. It emits a request; the guarding MCU checks it against deterministic rules before granting.

DISPENSE REQUEST →

Dispenser MCU checks: is this dose scheduled for now, within tolerance? Already dispensed (dose counter)? Cartridge matches the schedule? Only then does it unlock *exactly one* dose — and verify physically that one, and only one, left the mechanism.

MOTION REQUEST →

Supervisor MCU checks: is the commanded pose within joint limits? Does the trajectory stay under the human-safe torque/velocity ceiling? Is the collision sensor clear? Only then does it pass the command — while monitoring current draw in real time.

ESCALATION REQUEST →

Escalation logic checks: do the anomaly criteria hold, or is this a resident-initiated call? It places the call over an *independent cellular path* and does not depend on the cortex staying alive to complete it.

WHY THREE PROCESSORS, AND NOT ONE CLEVER ONE

The single-processor answer — “run the safety checks in a high-priority thread on the Jetson” — fails the exact test the founder set. A shared processor means a shared failure mode: a kernel panic, a memory-exhaustion event from the vision pipeline, or a bug in the VLA runtime can take the safety logic down with it.

Physical separation buys **independence of failure**. The supervisor and dispenser have their own power, clock, firmware, and watchdog on the cortex. When the smart part crashes — and over a fielded life it will — the dumb parts keep the arm still, the doses locked, and the phone line open.

CONNECTIVITY, DUAL-PATH BY DESIGN

Internal: CAN-FD between the MCUs and actuator subsystems — deterministic, differential, the automotive-grade safety-bus choice; the STM32H7 class supports FD-CAN and ECC memory natively.

External: home Wi-Fi for the bandwidth-hungry telepresence stream, plus an **independent cellular modem** dedicated to escalation so a downed router cannot silence a 911 call — the design MedMinder and MedReady already ship.

WHAT THIS ARCHITECTURE DELIBERATELY DOES NOT DO

It does not let the arm handle loose pills — that is the contact-rich, high-consequence task where learned manipulation is least reliable; the medication path is a sealed, mechanically-metered subsystem instead. It does not make the robot the diagnostician — vitals are measured, thresholded against clinician-set bounds, and escalated. And it does not put the emergency call behind the smart stack — escalation is deliberately the least intelligent, most independent path in the whole system.

THE LEARNED BRAIN — A $\pi 0$ -CLASS VLA

$\pi 0$ (Physical Intelligence) is a **PaliGemma ~3B vision-language backbone plus a ~300M flow-matching action expert** — ~3.3B parameters, 50 Hz control over ~1-second horizons. Open-weight base checkpoints are pre-trained on **10,000+ hours** of robot data. It is a *generalist* that fine-tunes on modest data (hours to ~20 hours per behaviour), and the published $\pi 0$ -ALOHA-towel checkpoint folds diverse towels zero-shot — a real fabric result.

Inference needs >8 GB GPU (RTX 4090 class); LoRA fine-tune >22.5 GB, full >70 GB (A100/H100). Inference is on-robot; training is off-robot. This requirement is why the on-robot compute is a Thor, not a smaller module.

THE PROPERTY THAT MUST NEVER BE FORGOTTEN

A VLA is a learned, probabilistic policy. Its authors say so plainly — the openpi README states outright that “ *$\pi 0$ may or may not work for you.*” It generalises impressively and it also fails, sometimes confidently.

That is acceptable for *bringing a cup* and unacceptable for *releasing a pill*. The architecture's entire job is to make sure the VLA is only ever wired to the first kind of task.

WHAT IS LEARNED VS. WHAT IS HARD-CODED — THE HONEST LEDGER

FUNCTION	WHERE IT RUNS	LEARNED OR DETERMINISTIC	FAILURE TOLERANCE
Recognise resident, detect a fall	Jetson (perception)	Learned (CV)	Tuned FP/FN; verified vs. contact sensors before escalation
Navigate, avoid obstacles	Jetson (Nav2 + SLAM)	Mixed	Graceful stop on uncertainty
“Bring me X,” grasp an object	Jetson ($\pi 0$ VLA)	Learned	May fail/retry; <i>never on the med path</i>
Estimate HR / respiration	Jetson (rPPG)	Learned	Characterised error; confirmed by contact sensor when it matters
Enforce torque / velocity / joint limits	STM32 supervisor	Deterministic	Hard stop; no drift possible
Collision e-stop	STM32 supervisor	Deterministic	Hardware-level
Watchdog on the Jetson	STM32 supervisor	Deterministic	Halts arm if cortex misbehaves
Release exactly one authorised dose	Dispenser MCU	Deterministic	Mechanically incapable of double-dose
Verify what was dispensed	Dispenser MCU	Deterministic	Physical confirmation, logged
Place the 911 call	Dispenser/escalation MCU + cellular	Deterministic	Survives cortex failure and router failure

The single most important row in this table is that no life-critical action appears in the “learned” column. That is the engineering content of “it can't drift.”

NOT SPECULATIVE

Diligent's Moxi — the most-deployed nurse-support robot in US hospitals, 1.25M+ autonomous deliveries across 25+ facilities — runs on the NVIDIA Jetson platform and trains in Isaac Sim. A fielded medical robot already uses this lineage.

Correctness is mechanical, not statistical.

A Geneva drive advancing sealed units *cannot* double-dose; a neural network deciding how many loose pills to pick can, and will, eventually. A good safe looks like a safe — the mechanism's security is meant to be visible and reassuring.

08 ACTUATION

The arm never touches medication. Dispensing is a sealed, mechanically-metered subsystem with its own controller.

THE MECHANISM, REASONED FROM THE REQUIREMENT OUT

- 1 **Pre-packaged unit doses**, not loose pills — loaded by pharmacy or caregiver. The machine advances sealed units; it never counts pills. Removes the whole class of “wrong count” errors at the source.
- 2 **Single-actuator hard-index drive** (Geneva / detented ratchet) — one command physically cannot advance two units.
- 3 **Physical lock between doses** — the entire supply is inaccessible except in an authorised window. The dementia-safe MedMinder-Jon property, in a robot.
- 4 **Deterministic authorisation** — released only when the schedule authorises, the dose is not already dispensed (monotonic counter the cortex cannot roll back), and the cartridge matches.
- 5 **Physical verification** — an optical gate plus a load-cell confirm one, and only one, unit left. A mismatch is logged as a fault, not marked taken, and the caregiver is alerted.
- 6 **Fail-locked on power loss** — medication secured, with a documented caregiver key override.

THE HONEST LIMITATION

The device confirms a dose was *dispensed and collected*; it cannot confirm the resident swallowed it — a limit shared by every dispenser on the market. The camera can add a supervised-swallow check as a soft signal; the honest claim is “dispensed and collected,” escalating on a missed collection.

MOBILITY & ARM, IN ONE LINE EACH

Base: differential drive, deliberately slow — the reliable, serviceable choice Moxi already proves. Prototype on ODrive S1 (\$149) + BLDC hub motors; every drive command passes the supervisor's hard speed ceiling.

Arm: light fetch-and-carry only. Prototype on Dynamixel-based research arm (Trossen WidowX class, ~\$3.1–4.6k); the servos' own current feedback gives human-safe collision detection — the quote-only ATI force sensor is a research instrument, designed out of the shipping product.

Contactless screens continuously and effortlessly; contact confirms on demand when a number must be trusted.

A technical reader will not accept “the camera reads your blood pressure” without the caveats, and neither should the founder. Getting this split right is the difference between a credible medical device and a gadget that cries wolf.

CONTACTLESS — THE ALWAYS-ON SCREENING LAYER

Trust it **Heart rate** — rPPG pooled bias ~-0.13 bpm vs. ECG; shipping software agrees within ~1-3 bpm front-facing, well-lit. The one vital the camera can be trusted on.

Trend only **Respiration** — camera RR error ~2.3 br/min; good for trend and anomaly, not a precise clinical number. And thermal: skin-temperature *indicator*, not core temperature.

Do not claim **Blood pressure & SpO₂** — contactless BP degrades without cuff calibration and is validated mostly on normotensive subjects. At best it flags category after calibration; it cannot replace a cuff.

Known failure conditions, stated: accuracy drops in low light and at elevated heart rates; motion and poor perfusion degrade the signal; skin-tone robustness must be tested, not assumed.

CONTACT — WHERE A NUMBER MUST BE TRUSTED, THE RESIDENT PARTICIPATES

Blood pressure Oscillometric upper-arm cuff — the only home-validated BP method. Robot prompts, times, logs.

SpO₂ Medical-grade OEM pulse-ox (Masimo/Nonin class) — far more reliable than any camera oxygen estimate.

Temperature Contact / clinical IR thermometer when the thermal-camera trend flags something.

TELEPRESENCE — THE EMOTIONAL CORE

An RGB-D camera (the robot's “eyes”) doubles as the caregiver window, streaming an encrypted two-way call. The feature that turns a monitoring device into a *presence*.

SUPPLY-CONTINUITY FLAG, VERIFIED

Intel is winding down the RealSense depth-camera business (select stereo lines continue; LiDAR/tracking EOL). The primary reference (D435i ~\$350 / D455 ~\$400) needs a qualified second source from day one — Orbbec (Gemini 335 ~\$250), Luxonis OAK-D, StereoLabs ZED. Moxi itself used a RealSense sensor, which underscores both that the choice is proven and that the whole industry now shares this second-sourcing problem.

FALL & ANOMALY DETECTION — CORROBORATION, NOT A BETTER CLASSIFIER

Fall detection fuses depth pose, thermal, and audio — but it is a *learned* detector, so it does not escalate on its own certainty alone. A suspected fall triggers the robot to approach, attempt verbal contact, offer the vitals check, and open the caregiver window; escalation to 911 follows deterministic criteria, not a single model's guess. The two-path logic keeps false alarms low without risking a missed emergency.

Two failure modes can kill. The architecture mitigates each one *independently* — no single fault, including a fault in the AI, can cause the harm.

WRITTEN AGAINST THE RECOGNISED STANDARDS

- ISO 14971
- IEC 62304
- IEC 60601-1
- IEC 62366-1
- ISO 13482*

*ISO 13482 is expressly not for medical applications — used as a reference for the robot-specific hazards (contact, mobility) alongside the medical-device standards.

Hazard 1 — Medication error

wrong drug · wrong dose · double dose · wrong time

CAUSE	CONTROL, AND WHERE IT LIVES	WHY INDEPENDENT OF THE AI
Double dose	Mechanical hard-index dispenser (Geneva/detent) advancing exactly one sealed unit	Physics; no software can override it
Wrong time / repeated	Deterministic schedule check + monotonic dose counter on the dispenser MCU	Separate processor; the cortex cannot roll back the counter
Wrong drug	Sealed unit-dose packaging loaded by pharmacy; cartridge-to-schedule match	Drug identity fixed at load time by a human, outside the device's decision
AI “decides” to dispense	The VLA has no dispense authority — it can only request; the MCU grants against rules	By construction: the learned model is not wired to the release actuator
Released but not delivered (jam)	Optical-gate + load-cell verification ; mismatch → logged fault, not marked taken	Physical sensor threshold on the MCU; audible

Residual risk, stated honestly: a loading error by the human/pharmacy, and non-ingestion after collection — the risks the whole industry carries, far smaller than “an AI decided how many pills to pick up.”

Hazard 2 — Missed emergency escalation

minimise misses without so many false alarms the device gets switched off

CAUSE	CONTROL, AND WHERE IT LIVES	WHY INDEPENDENT
Detection model misses the event	Multi-signal corroboration + a resident panic path (button/voice) needing no model	The manual path bypasses the AI entirely
The smart stack has crashed	Escalation trigger & dialing live on the guarded MCU , which watchdogs the cortex	Physically separate processor, own power
Home network / router down	Independent cellular modem dedicated to escalation	Does not share the Wi-Fi telepresence path
Device loses power	Battery backup sized to keep escalation + dispenser lock alive through an outage	Dedicated power budget for safety functions
False-alarm fatigue → ignored	Graduated escalation: approach, engage, offer vitals, open window — <i>then</i> 911 on deterministic criteria	911 is the corroborated last step, keeping true positives trusted

The escalation decision is deliberately the least intelligent path. 911 is placed on deterministic criteria — a panic press; or no response to contact within a window plus a corroborating signal; or a vitals reading crossing a clinician-set hard threshold — biased toward escalation, because a false 911 call is recoverable and a missed one is not. The honest claim is a characterised, corroborated system with an always-available manual path and independent power and connectivity — not omniscience.

Built to survive a technical reader who checks the numbers.

A direct-materials BOM: it excludes NRE, tooling, labour, assembly, test, certification and margin — which are larger than the parts and are the finance section's job. Verified prices carry a source; ranges are honest engineering estimates, labelled.

SUBSYSTEM	REFERENCE PART / CLASS	PROTOTYPE (1)	VOLUME (10K/YR)
High-level compute	NVIDIA Jetson Thor long lead	\$5,499	\$3,000–4,000
Safety supervisor MCU	STM32H7 dual-core + board	\$50–120	\$8–15
Dispenser MCU	STM32H7/G4 class + board	\$50–100	\$6–12
Manipulation arm	Proto: Trossen WidowX; Vol: custom 4–5-DOF	\$3,100–4,600	\$700–1,400
Wrist F/T (optional)	ATI Nano17 / Mini45 quote-only	\$3,000–7,000	designed out
Mobile base — drive	Proto: ODrive S1x2 + BLDC; Vol: custom hub	\$600–1,000	\$250–450
Chassis, wheels, structure	Machined → molded	\$800–1,500	\$300–600
Depth cameras (x2)	RealSense; 2nd-source Orbbec/ZED supply risk	\$700–900	\$400–700
2D LiDAR · thermal	Slamtec RPLIDAR · FLIR Lepton class	\$300–650	\$230–500
Contact vitals (BP/SpO ₂ /temp)	Home-validated cuff · OEM pulse-ox · IR	\$210–560	\$115–345
Med-dispensing mechanism	Proto: machined carousel; Vol: injection-molded	\$500–1,200	\$150–400
Display · audio · connectivity	Touchscreen · mic array + speaker · Wi-Fi 6E + cellular	\$310–800	\$135–370
Battery + power · wiring & misc	Li-ion pack + BMS + fail-safe backup · CAN-FD bus	\$700–1,600	\$350–800
Headline band carried across this dossier		\$28,000–42,000	\$4,200–6,800

RECONCILING THE HEADLINE

The **prototype \$28,000–42,000** band is a realistic *engineering* unit: the base parts sum is lower, but the headline bakes in the research-grade arm, the optional ATI sensor for the research phase, spares, and the redundant second-of-everything a bring-up demands. Parts only — the true first-unit program cost (firmware, mechanical design, safety engineering, iteration) is far larger.

THE VOLUME LEVERS & THE ±30% BAND

The **volume \$4,200–6,800/unit** is a realistic post-DFM band — tighter than a naive sum-of-highs. Three levers move it most: compute (the largest, least-negotiable line), the arm (research → purpose-built is the biggest single saving), and the injection-molded dispenser. It assumes DFM work not yet done and that the compute requirement holds — a ±30% band is the honest posture at this stage. *A live workbook companion lets the client flex volume, sourcing, and the compute assumption and watch the per-unit number move.*

A build story a robotics engineer will believe, because it matches what the field actually delivers.

You cannot safely or affordably collect the training data for a home nurse robot on real elderly people. Simulation generates the volume and diversity, red-teams dangerous edge cases without anyone getting hurt, and validates a new model before it touches hardware. This is established practice: **Diligent's Moxi is trained and tested in NVIDIA Isaac Sim and Isaac Lab** — a fielded medical robot already lives this workflow.

ISAAC SIM / LAB

The sensor-faithful, visual twin: photorealistic RTX, GPU-parallel (~82k–94k FPS across ~4,096 envs), physics + rendering domain randomisation. Closes the *visual gap* for the vision-heavy perception and navigation stack.

MUJOCO / MJX / NEWTON

The contact & control twin: fast, accurate contact dynamics (~2.7M steps/sec on TPU). Where the manipulation control is trained and stressed — contact is where physics simulators are most tested.

THE RECIPE, STATED PLAINLY

Pretrain broad skills in simulation with heavy domain randomisation, then fine-tune on limited real-world data.

The documented standard, not an invention. The claim the dossier makes is precise: sim gives safe, scalable coverage and closes the gap on navigation and rigid tasks; the deterministic safety mechanisms are validated by conventional V&V, not simulation; and contact-rich and soft-object skills are honestly scoped as real-data efforts.

WHAT TRANSFERS WELL — AND WHAT DOES NOT (THE HONEST GAP)

Transfer success correlates inversely with contact complexity. Free-space motion and simple rigid contacts transfer well; contact-rich manipulation and deformable/soft objects transfer poorly regardless of technique.

CAPABILITY	PROGNOSIS	CONSEQUENCE FOR THE PROGRAM
Navigation, obstacle avoidance	Transfers well	Train primarily in sim; modest real-world validation (84–93% with tuned DR)
Person / fall detection	Transfers reasonably	Sim for coverage, real data to calibrate the operating point
Rigid-object fetch-and-carry	With fine-tuning	Sim pretrain + real demonstrations per object class
Contact-rich manipulation	Transfers poorly	Budget real-robot data collection; do not promise sim-only
Fabric / soft-object handling	The honest weak spot	Real data essential; scope conservatively; keep first release to rigid fetch-and-carry

THE MED-DISPENSER IS DELIBERATELY OUTSIDE THIS GAP

Note what is *not* in the table: releasing medication. That is precisely because the dispenser was designed as a deterministic mechanism rather than a learned manipulation task — it sidesteps the entire contact-rich sim-to-real problem. This is the architecture paying off: the one task that cannot tolerate the sim-to-real gap is the one task that never depends on a learned policy in the first place.



PART THREE — THE MOAT

The hardest problem is the moat. A mapped regulatory path is what turns a scary chart into a fundable one.

FDA classification, the pre-submission that de-risks the raise, the AI/ML framework that rewards the “can't drift” architecture, HIPAA, liability, and the phased path with the real fees and clocks. Every regulatory fact verified against FDA.gov, the eCFR, or HHS.gov.

Prepared as strategy for the company and its regulatory counsel — not a regulatory determination. Classification, predicate and pathway are confirmed with FDA via the Pre-Submission and a qualified regulatory consultant.

A system of functions that classify separately — and the classification of the whole is driven by the highest-risk function inside it.

A serious reader who has done medical-device work reaches for this distinction immediately. The correct exercise is to find the product code that matches each function's intended use — not to negotiate a class in the abstract.

FUNCTION	LIKELY PRODUCT CODE / REGULATION	CLASS	BASIS
Scheduled Rx dispensing (dependent patient)	NZH · 21 CFR 880.6315 Remote Medication Management System	Class II special controls	FDA special-controls guidance, 2007 — a direct match to intended use
Multi-parameter vitals, home, Rx	DRG/FLL · 870.2910 or MHX · 870.1025	Class II	Living predicates incl. Vitls K191620 (home, Rx)
Caregiver camera telepresence	DRG · 21 CFR 870.2910	Class II	Predicate RP-VITA K123229
Anomaly detection → autonomous 911	Novel SaMD — see next spread	Class II De Novo probable	No clean predicate for autonomous escalation

THE CLASS I TRAP

FDA has a Class I, 510(k)-exempt code for a solid medication dispenser (NXB, 21 CFR 890.5050). If the device were just a pill organiser that pops open a tray on schedule, it could plausibly sit there and need no premarket submission. **This is the reading a founder wants to be true, and it is the reading that gets a company a warning letter.** The Class I “daily activity assist” codes cover devices that *aid* a person still managing their own regimen — not a device that autonomously selects, verifies, and delivers a prescribed dose to a dependent patient.

THE GIFT INSIDE THE CORRECT READING

FDA's NZH guidance calls the dispensing software a “**major level of concern**” for exactly the founder's fear — “a patient could miss necessary medication, or receive an incorrect dose or the wrong drug... if the device fails because of a software defect.” It then prescribes the mitigations FDA expects, which map one-to-one onto the engineering safety case: device-side medication identification, verify correct dose delivered, verify inventory, physical segregation, complete event log.

The company is not guessing at “safe enough.” **FDA has published the acceptance criteria**, and the deterministic dispensing subsystem is being designed to meet a documented federal standard, not an internal opinion.

Bottom line for an investor: this is a **Class II medical device** — 510(k) for the functions that have predicates, very probably a De Novo for the novel autonomous-escalation function. It is *not* a Class III / PMA device. The distinction is worth stating plainly to a nervous reader: Class II 510(k)/De Novo is a navigable, precedent, ~~sub-\$75K in fees~~, roughly-a-year-per-submission path — not the PMA nightmare they may be picturing. And the RP-VITA precedent converts “a robot has never been cleared to do this” (false) into “here is the K-number of the one that was” (true).

A free, voluntary FDA meeting that converts regulatory risk from unpriced to priced.

The single most fundable regulatory event this company can execute before it has a cleared product is a Pre-Submission ("Pre-Sub") meeting with FDA — part of the Q-Submission Program, governed by FDA's guidance finalised **May 29, 2025**. It lets the company put specific questions to the actual review division *before* filing any marketing submission, and get the Agency's position in writing.

A pre-seed investor's deepest fear on a device like this is regulatory ambiguity — that FDA comes back a year and \$30M later and says "this is a harder pathway than you thought." A Pre-Sub, done well, retires that specific risk. Walking out with FDA's written agreement on classification, predicate, and testing strategy converts the story from "*we believe this is Class II 510(k)/De Novo*" to "**FDA has told us, in writing, this is Class II 510(k)/De Novo, and here is the letter.**"

That letter is a Series-A asset — the single document that most cheaply moves the company's regulatory risk from unpriced to priced.

THE RECOMMENDED SEQUENCE

- 1 Pre-seed** funds the prototype + regulatory groundwork: engage a consultant, draft the intended-use statement and predicate analysis, start the ISO 14971 / IEC 62304 files.
- 2 File the Pre-Sub** on the phased product (monitoring/telepresence first, cleanest predicates) with a few sharp questions — FDA caps a Pre-Sub at four topic areas.
- 3 The FDA meeting & written feedback is the pre-seed's headline deliverable** and the Series-A trigger.
- 4 Series A** funds the clinical/human-factors and bench testing FDA asked for, and the actual submissions.

\$0

USER FEE FOR A PRE-SUBMISSION

It is free to file; the cost is the effort to write a package good enough to get useful answers. Timeline: ~15-day acceptance review, then written feedback by ~day 70 — a **60–90-day cycle**.

CLEARED PREDICATES — REAL K-NUMBERS

K123229 Class II

RP-VITA Remote Presence Robot

(InTouch/iRobot), cleared 11/20/2012 — the first autonomous-navigation remote-presence robot cleared by FDA.

K191620 Class II

Vitals Platform, cleared 06/01/2020 —

wireless multi-parameter vitals, Rx use, "Home and Healthcare Settings." Proves home multi-parameter monitoring is an established, clearable Class II claim with living predicates.

The founder's constraint is, almost word for word, the central concern FDA has spent years building a framework around.

IS THE ESCALATION SOFTWARE REGULATED? YES — AND OWNING THAT IS A STRENGTH

The 21st Century Cures Act carved certain Clinical Decision Support software *out* of the device definition. The escalation function fails the carve-out by design, on two of the four criteria:

Fails Criterion 1 (data source) — it continuously acquires and analyses physiological signals, exactly the signal-pattern analysis the carve-out excludes.

Fails Criterion 3/4 (support vs. replace) — a device that autonomously calls 911 is not enabling a clinician to independently review before acting; it *is* the action, and the user is a patient/caregiver, not a reviewing clinician.

The conclusion, stated as a strength: **the anomaly-detection-and-escalation software is regulated SaMD**. Any pitch that claims it slips through the carve-out is wrong in a way FDA and a diligence lawyer both catch. There is no clean 510(k) predicate for autonomous emergency escalation — which points to a **De Novo**, and that is good news: a granted De Novo creates a new product code the company *owns* as the reference point competitors must later cite. Being the device that *defines* the category is a durable regulatory moat.

DE NOVO, BY THE CLOCK

319 days

Median from receipt for grants 2023–2025; FDA's FY2027 goal is 150 FDA days for 90%. Plan ~10–16 months.

THE RECOMMENDED FRAMING

Treat the novel escalation function as a probable De Novo — and treat that as a category-defining regulatory asset, not just a clearance. The Pre-Sub is exactly how the company gets FDA to pick the route.

THE PCCP — WHERE THE ENGINEERING AND THE REGULATORY FRAMEWORK BECOME ONE SENTENCE

FDA's worry about AI/ML is precisely *drift*: a model that changes behaviour after clearance is a model whose safety was never reviewed. FDA's answer is the **Predetermined Change Control Plan (PCCP)** — finalised December 3, 2024 — documentation in the original submission that pre-specifies exactly what modifications the AI may undergo, how they'll be validated, and bounds them to the intended use. In-scope changes ship without a new submission; anything outside still requires one.

The safety-critical decisions live in locked, deterministic, verifiable logic — not in an adaptive model. The AI never has final authority over a pill or a 911 call. In FDA/IEC 62304 terms, the highest-harm software items are locked and held to the strictest verification — exactly the “can't drift” property FDA wants. Where ML legitimately improves (refining anomaly sensitivity), a PCCP makes that improvement lawful and bounded.

The deterministic-safety-core architecture is not just good engineering — it is the specific design FDA's AI framework is built to reward, and it is the same design that answers the founder's “can't drift” mandate. The engineering safety case and the regulatory strategy are, by construction, one story.

HIPAA: the camera stream and the vitals are PHI.

A breach is an existential brand event for an elder-care product, and the compliance architecture is not optional. The safe and fundable posture is to **architect for full HIPAA compliance from day one regardless of channel** — far cheaper to build in than retrofit, and enterprise/provider channels demand it in procurement.

WHERE THE COMPANY SITS

Sold through providers, agencies or payers — the likeliest channel — the company is almost certainly a **business associate**, signing a BAA with each, and directly liable under HITECH §13401. It must require BAAs from every downstream subprocessor that touches ePHI.

WHAT THE ARCHITECTURE MUST CARRY

Encryption in transit and at rest (OCR's January 2026 NPRM moves it from *addressable* to **mandatory** — build to the proposed standard now); access controls and minimum-necessary; audit logging of every stream access; patient authentication; contingency/incident response.

DIGNITY-BY-DESIGN IS THE DIFFERENTIATOR

Real, revocable consent that handles surrogate decision-making; a hardware indicator, physical shutter and audible cue so the resident always knows when the camera is live; privacy-sensitive zones and times handled deliberately; default toward *not* retaining continuous video — process locally, transmit on event, store minimally. A privacy win, a breach-surface reduction, and a cost reduction at once, and it shows favourably in FDA human-factors review.

Liability: the architecture is the defense.

A device that autonomously dispenses medication and decides whether to summon help sits in a high-exposure corner of product liability — design defect, manufacturing defect, failure to warn — and against a *standard of care*: “a competent human caregiver would have caught this.” It cannot be waved away; the honest answer is also the argument for the architecture.

WHY THE DETERMINISTIC CORE IS THE LIABILITY DEFENSE

A locked, deterministic safety core is auditable and reproducible. “The AI decided” is an indefensible sentence in a deposition; “the deterministic controller executed a verified, logged, FDA-reviewed decision rule” is a defensible one. The two-brain split is a legal architecture as much as an engineering one.

AUTONOMOUS 911 IS THE LAST RUNG OF A LADDER, NOT A REFLEX

- 1 On anomaly, alert & attempt to engage the patient
- 2 Escalate to the caregiver for confirmation within a bounded window
- 3 Only on caregiver non-response *and* a high-confidence, high-severity signal → autonomous 911

Every tier, threshold and decision is logged — the complete event log (required by FDA's NZH controls anyway) is the company's contemporaneous evidence in any dispute.

THE POSTURE, PLAINLY

The company does not eliminate liability — no medical device does. It *manages* it: humans in the loop, safety-critical decisions in verifiable logic, documented federal standards met, everything logged, adequate warnings, appropriate insurance. The difference between “assists a caregiver” and “replaces a caregiver” is worth real money, and is settled with counsel.

A navigable Class II route — not a PMA — with under ~\$75K in total agency fees and a free Pre-Sub that de-risks everything after it.

A technical reader knows roughly what clearance costs and takes, and will trust an honest set of numbers far more than optimistic ones. The figures below are FDA's own published fees and clocks — not the larger, more variable internal cost of building the submission, which is called out separately.

FDA USER FEES (MDUFA V, FY2027)

SUBMISSION	STANDARD	SMALL BUSINESS
510(k)	\$28,653	\$7,163
De Novo request	\$191,020	\$47,755
Pre-Submission (Q-Sub)	\$0	\$0
Est. registration (annual)	\$13,785	\$13,785

Small-business rate = 25% of standard for these submission types (gross receipts ≤ \$100M). The company must apply for Small Business Determination status *before* filing. Source: FDA MDUFA fee schedule FY2026/FY2027.

REVIEW CLOCKS (VERIFIED)

Pre-Submission	~60-90 days	to FDA written feedback · no fee
510(k)	~6-12 mo	112-day TTD goal; ~149-day recent average
De Novo	~10-16 mo	median 319 days; FY2027 goal 150 FDA days

WHAT EACH STAGE OF CAPITAL BUYS

- Pre-seed (~\$3.5M)**
The Phase-1 prototype, the regulatory groundwork, and the FDA Pre-Submission with written feedback in hand. It does *not* buy a clearance, and the pitch does not claim it does. It buys the milestone that makes the Series A rational.
- Series A**
The first clearance (Phase-1 monitoring/telepresence 510(k)) — roughly 12-18 months from a well-run Pre-Sub, at FDA fees under ~\$10K plus the six-figure cost of testing and submission prep.
- Series B & revenue**
The hard clearances — the De Novo escalation, the NZH dispensing — are later-stage, multi-year efforts, correctly sequenced to *last*.

THE ONE-SENTENCE VERSION FOR THE DECK

The FDA path for this device is a navigable Class II 510(k)/De Novo route — not a PMA — with under ~\$75K in total agency fees, a ~1-year review clock per submission, and a free Pre-Submission that lets us get FDA's written agreement on the whole strategy before we spend the first submission dollar; the pre-seed is sized to reach that Pre-Sub, which is the milestone that de-risks everything after it.

Every load-bearing fact, with the source it was verified against.

Figures marked RANGE are honest engineering estimates, not quotes, and are labelled as such in-line. Where a figure is properly a range, or reputable sources disagree, both are shown and the conservative one anchors the case.

COMPUTE & AUTONOMY

Jetson Thor (\$5,499 dev kit; T5000 \$4,999 / T4000 \$2,999 @1KU+; 128 GB, 2,070 FP4 TFLOPS, 40–130 W; allocation-constrained) — NVIDIA Jetson FAQ via Physical AI Lab (Aug 2026); VideoCardz (Jul 2026). The “\$3,499” in NVIDIA's Aug 2025 blog is stale.

STM32H7 (Cortex-M7 480–600 MHz, ECC on Flash/RAM, FD-CAN; ~\$8.19/unit @10k) — STMicroelectronics via Circuit Digest, EETimes, ST product pages, 2026.

π0 (PaliGemma ~3B + ~300M action expert, ~3.3B total, 50 Hz; open-weight, 10k+ hrs; inference >8 GB GPU; “may or may not work for you”; π0-ALOHA-towel 0-shot) — Physical Intelligence openpi repo + π0 paper arXiv:2410.24164, 2026.

ACTUATION & SENSING

Dynamixel X-series torque figures & XM430-W210-T \$310.39 — ROBOTIS documentation / ROBOTIS America, 2026. **Trossen WidowX** 250 S ~\$3,100 (discontinued), WidowX AI from \$4,545.95 — Trossen / Robotics Center, 2026. **ATI Nano17/Mini45** quote-only, built-to-order — ATI/Sensor360/Max-Robotics, 2026. **ODrive S1** \$149 &c. — ODrive Robotics US, 2026.

rPPG (HR ~1–3 bpm, pooled bias ~–0.13; RR ~2.3 br/min; BP unreliable without cuff) — Bautista et al. 2023; Pham et al. 2022;

MDPI Bioengineering 2026; Nature npj Digital Medicine 2025 (CHILL); OAE Conn. Health 2022. **RealSense** D435i ~\$350 / D455 ~\$400; Intel winding down; 2nd-source Orbbec/ZED — Intel/Mouser + CRN + The Robot Report + Tom's Hardware, 2025–2026. **Slamtec RPLIDAR A1/A2** — Slamtec datasheets, 2026.

DISPENSER MARKET REFERENCE

Hero (unlocked cup, key access on power fail, not dementia-safe) — Hero Instruction Manual 3.0 (Dec 2024) + reviews. **MedMinder Jon / MedReady / MedaCube** (hard-locking per-dose, built-in cellular; MedaCube FDA-registered; \$30–\$125/mo; 8–48 h battery) — product documentation + caregiver review sites, 2026.

SAFETY STANDARDS & SIM

ISO 14971 · IEC 62304 · IEC 60601-1 · IEC 62366-1 · ISO 13482 (the last expressly not for medical applications) — FDA recognized-standards database; IEC/ISO; Johner Institute; RESNA care-robot study, 2026. **Isaac Sim/Lab** (~82k–94k FPS @4,096 envs), **MuJoCo** (MJX ~2.7M steps/sec), **Newton**; contact-rich & deformable transfer poorly, 84–93% with tuned DR — Isaac Lab paper/docs; CodeSOTA survey; Robotics Center 2026; MLR proceedings. **Moxi** (Jetson + Isaac Sim, RealSense + Hokuyo, 25+ hospitals, 1.25M+ deliveries) — robotsguide.com; Serve/Diligent release, 2026.

DEMOGRAPHICS, COST & WORKFORCE

Population Reference Bureau, *Aging in the United States* (Census 2023 projections); US Census P25-1144; Administration for Community Living, *2023 Profile of Older Americans*; Peterson Foundation / US 2050. Care costs: Genworth / CareScout *2024 Cost of Care Survey*. Workforce: PHI *Direct Care Workforce*; BLS Employment Projections 2024–2034; Commonwealth Fund 2024; HRSA. Caregiving value: AARP, *Valuing the Invaluable* (2023) — \$600B, 36 billion hours.

ADHERENCE & MARKET SIZING

Non-adherence \$100–290B — *Annals of Internal Medicine* 2012 systematic review. Non-optimized medication therapy \$528.4B (2016), ~275,000 deaths — Watanabe, McInnis & Hirsch, *Annals of Pharmacotherapy* 2018 (kept distinct from the non-adherence figure). Market: MarketsandMarkets (RPM \$27.7B→\$56.9B, the conservative anchor); Grand View (home healthcare); Strategic Market Research (smart dispenser \$2.1B→\$4.2B); alternative RPM estimates from Global Insight Services and einpresswire noted for range, not blended.

REIMBURSEMENT

CMS CY2024 & CY2025 Physician Fee Schedule Final Rule (conversion factor \$32.35), as summarized by Smart Meter, Health Recovery Solutions, Tenovi, PYA — RPM CPT codes 99453/99454/99457/99458/99091.

COMPETITORS & FUNDING

Labrador Systems (\$3.1M seed, Jan 2022, Alexa Fund & iRobot Ventures) — PRNewswire. Diligent Robotics (\$50M+ total; ~\$50M through 2022 Series B + a \$25M round reported 2023) —

MobiHealthNews, TechFundingNews. Intuition Robotics (~\$83M; \$25M 2023, Woven Capital) — TechCrunch, OurCrowd. Hero Health (\$12M Series A; pricing) — MedCity News, Pulse 2.0, herohealth.com. MedMinder, Sensi.AI, CarePredict, Aloe Care — industry funding trackers & Crunchbase-sourced summaries.

FDA — CLASSIFICATION & PREDICATES

Product codes NXB (890.5050, Class I exempt), KYX (880.6430), NZH (880.6315, Class II special controls — guidance issued Oct 19, 2007), MHX (870.1025), DRG (870.2910) — FDA Product Classification database / eCFR. Cleared predicates: **K123229** RP-VITA (cleared 11/20/2012); **K191620** Vitis Platform (cleared 06/01/2020) — accessdata.fda.gov 510(k) records.

FDA — PATHWAYS, AI/ML, FEES

Q-Submission Program guidance (final, May 29, 2025); De Novo pathway (median 319 days, grants 2023–2025); PCCP guidance (final, Dec 3, 2024, FD&C Act §515C); Clinical Decision Support Software guidance (2026 final; Cures Act §3060(a); FD&C §520(o)(1)(E)). MDUFA FY2026/FY2027 fee schedule (510(k) \$28,653 / \$7,163; De Novo \$191,020 / \$47,755; registration \$13,785); MDUFA Performance Goals FYs 2023–2027 (112-day TTD goal; ~149-day recent average).

HHS — HIPAA

HHS *Summary of the HIPAA Security Rule* (45 CFR Part 160 & 164); covered entities & business associates (160.103); HITECH §13401 (direct BA liability); BAA requirements (164.308(b)); required vs. addressable specs (164.306(d)); HHS OCR telehealth guidance & the January 2026 Security Rule NPRM (encryption/MFA toward mandatory).

A NOTE ON THE STANDARD HELD THROUGHOUT

“Verified” here means read on the source’s own page during preparation of this dossier. Figures labelled “assumption” in the market model are disclosed as modelling inputs, not facts. The regulatory analysis is prepared as strategy for the company and its counsel — not a regulatory determination; classification, predicate and pathway are confirmed with FDA via the Pre-Submission and a qualified regulatory consultant. A market size is a claim with a source or it is a guess — this document labels which throughout.

The nurse that doesn't leave at the end of a shift — engineered so the parts that must never be wrong are the parts that cannot drift.

\$3.5M

PRE-SEED, MILESTONE-
GATED

Class II

510(K)/DE NOVO, NOT
PMA

Pre-Sub

THE MILESTONE THAT DE-RISKS
THE A