nightingale os

trauma informed. Nurse forward. Systems that care.

MARKET ANALYSIS — FOR INVESTORS

The Unsustainable System

U.S. healthcare economics, the failure of digitization, and the operational-control opportunity

Regulation Loop, Inc. d/b/a Nightingale OS
Delaware C-Corp · Pre-Seed · Prepared for investor diligence
June 2026

About this analysis. Sections I-IV are an objective analysis of the U.S. healthcare system built entirely on public data from government and peer-reviewed sources (CMS, OECD, CDC, the Commonwealth Fund, JAMA/NEJM, ACEP, AHA, and others), cited inline and listed in full at the end. Section V presents market positioning and is the authors' interpretation. Forward-looking statements and any figures drawn from the Nightingale OS digital pilot are modeled projections, not observed outcomes, and are labeled as such.

Contents

Executive summary

The United States operates the most expensive healthcare system in human history and one of the worst-performing in the developed world. In 2023 it spent $4.9 trillion, or 17.6% of GDP, and the federal actuary expects that share to keep climbing toward 20% (CMS, 2024). For that money the country buys a life expectancy 3.7 years below comparable nations, the highest maternal and infant mortality in the high-income world, and a health system the Commonwealth Fund ranked dead last among ten wealthy peers in 2024. The system is not under-resourced. It is mis-operated.

The central finding of this analysis is that roughly $760-935 billion a year, about one in four dollars, is waste (Shrank et al., JAMA 2019), and its largest and most addressable categories are **operational and administrative, not clinical**. The failure is one of coordination, flow, staffing, and documentation: the logistics of care. Nurses turn over at 17.6% a year at a replacement cost of $60,090 each (NSI, 2026). Emergency departments board patients in hallways in what the nation's emergency physicians call a public health crisis (ACEP). About a quarter of inpatient days are avoidable (Advisory Board), and one in four hospital admissions involves an adverse event (Bates et al., NEJM 2023).

The industry's answer for two decades was technology, and it has not worked. The federal government spent more than $35 billion through the HITECH Act to drive electronic health records from 3% to 96% hospital adoption (Health Affairs, 2016). Costs rose, burnout rose, and outcomes did not. The reason is structural: **the EHR digitized the analog record without re-architecting the operation**. It is a system of record, not a system of action. The wave of point solutions that followed (command centers, ambient scribes, scheduling tools) each optimizes one local process while the system as a whole stays uncoordinated.

> **The thesis.** Healthcare does not need a new idea. It needs an old one that every other complex, high-stakes domain already adopted: operational and logistical control over the entire system. The military perfected it. Operations research was born in World War II to solve exactly this class of problem, and its core mathematics, queuing theory, explains why hospitals fail. Hospitals try to optimize the productivity of individual humans, running at ever-higher utilization. Queuing theory proves that as utilization approaches capacity, delay does not rise gradually; it explodes. Running hot works until it doesn't, and it isn't working. The alternative is to automate the analog coordination work that consumes clinicians and to manage the system for flow. Automating the analog unlocks the most precious and most wasted resource in the building: human beings.

For investors, the opportunity is framed by a simple gap. The hospital capacity-management software market is about $3.8 billion a year (Grand View, 2024), and all U.S. digital-health venture funding was $10.1 billion in 2024 (Rock Health), against roughly $850 billion in annual operational waste those tools are meant to address. The category that wins will not be a better dashboard bolted onto the record. It will be the **system of action** that sits above the records, coordinates the whole operation, and is auditable by design. Nightingale OS is built to be that layer. The rest of

this document makes the objective case for the problem (Sections I-IV), then the market case for the opportunity (Section V).

I. An unsustainable system: the macroeconomic case

I.1 The cost curve

U.S. national health expenditure reached $4.9 trillion in 2023, growing 7.5% in a single year and consuming 17.6% of gross domestic product, up from 17.3% in 2022 (CMS Office of the Actuary, 2024). The trajectory is the story. Health spending was about 5% of GDP in 1960; it has roughly tripled as a share of the economy over six decades, and the federal actuary projects it will reach 19.7% of GDP by 2032 as health spending grows 5.6% annually against GDP growth of 4.3% (CMS, 2024). The Commonwealth Fund projects the share will exceed 20% by 2035.

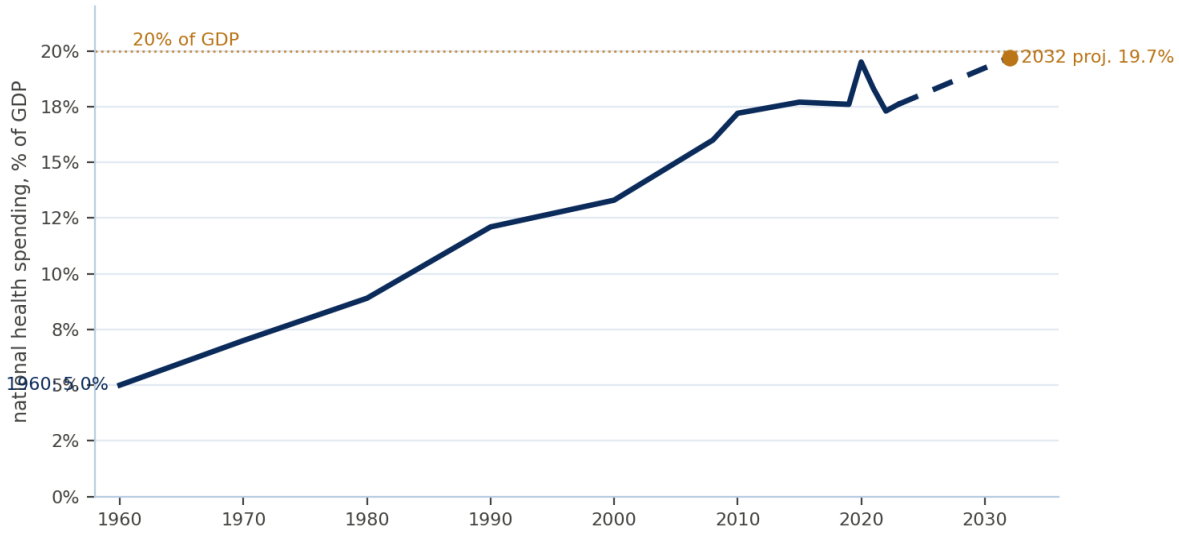

Figure 1. National health spending as a share of GDP, 1960-2023 (actual) with the CMS projection to 2032. The 2020 spike reflects both pandemic spending and a contracting denominator. Source: CMS Office of the Actuary.

A cost line that compounds faster than the economy funding it is, by definition, unsustainable. It must eventually consume an impossible share of output, crowd out every other public and private priority, or break. The question for a serious analyst is not whether the curve bends but how, and what drives it. Americans do not consume dramatically more care than their peers; the gap is higher prices and higher administrative and operational overhead per unit of care.

I.2 The value mismatch: most spending, worst outcomes

The defining feature of U.S. healthcare is the gap between what it costs and what it delivers. The United States spends roughly $14,775 per person per year, about 88% more than the ~$7,860 average of comparable wealthy nations, yet ranks at or near the bottom of that peer group on nearly every outcome that matters (Peterson-KFF, 2025).

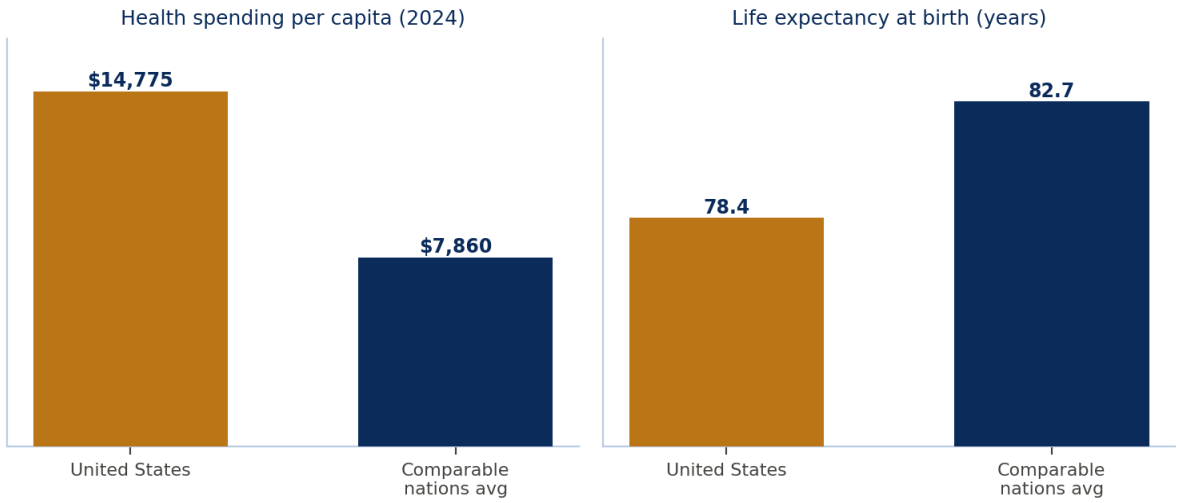

Figure 2. The United States spends far more per person yet lives several years less than the average of comparable high-income nations. Sources: Peterson-KFF Health System Tracker (spending, 2024); CDC/NCHS and Peterson-KFF (life expectancy).

The specifics are stark. U.S. life expectancy was 79.0 years in 2024, a recovery from the pandemic trough but still 3.7 years below the comparable-country average of 82.7 (CDC/NCHS; Peterson-KFF). Maternal mortality, at 17.9 deaths per 100,000 live births, is the highest in the high-income world, and more than three times higher for Black women (CDC, 2026). Infant mortality ranks 32nd of 38 OECD nations (OECD). Avoidable mortality runs about 50% above the OECD average and **rose in every U.S. state** between 2009 and 2019 even as it fell across peer countries (OECD; JAMA, 2025). In its 2024 Mirror, Mirror international comparison, the Commonwealth Fund ranked the U.S. health system last overall among ten high-income countries, and last specifically on access to care and on health outcomes.

Beneath the outcomes sits a chronic-disease burden the system manages poorly: 42% of U.S. adults have two or more chronic conditions, adult obesity has reached 40.3%, and roughly 38% of Americans, including the insured, reported skipping or delaying needed care because of cost in 2024 (CDC; Commonwealth Fund). An estimated 100 million people carry medical debt totaling at least $220 billion (KFF; Peterson-KFF). A system can be expensive or it can ration; the United States manages both at once.

I.3 The fiscal cliff

The unsustainability is not abstract; it has a date attached. The Medicare Hospital Insurance trust fund, which pays for inpatient care for 68 million Americans, is now projected to be depleted in the second quarter of 2033, at which point incoming revenue would cover only 89% of scheduled benefits (Medicare Trustees Report, 2025). That depletion date moved **three years closer** in a single year's report, driven by higher-than-expected hospital, hospice, and drug spending. Medicare and Medicaid together are the largest drivers of long-term federal debt projections, and hospitals are simultaneously squeezed from the payer side: Medicare reimbursed roughly 83 cents per dollar of care delivered in 2024 (AHA).

The burden lands on households and employers as well as the Treasury. The cost of family health coverage has risen faster than wages for two decades, consuming a growing share of total worker

compensation and suppressing take-home pay. The political economy of a sector approaching a fifth of GDP, with a trust fund on a countdown clock and households already rationing care, is the definition of a system under terminal strain. Reform is not optional; it is arithmetic.

I.4 Waste, not scarcity

The most important fact for an investor is where the money goes. The landmark estimate, Shrank, Rogstad, and Parekh in JAMA (2019), put waste in U.S. healthcare at **$760-935 billion per year, roughly 25% of all health spending**. The largest categories are administrative and operational rather than clinical overtreatment.

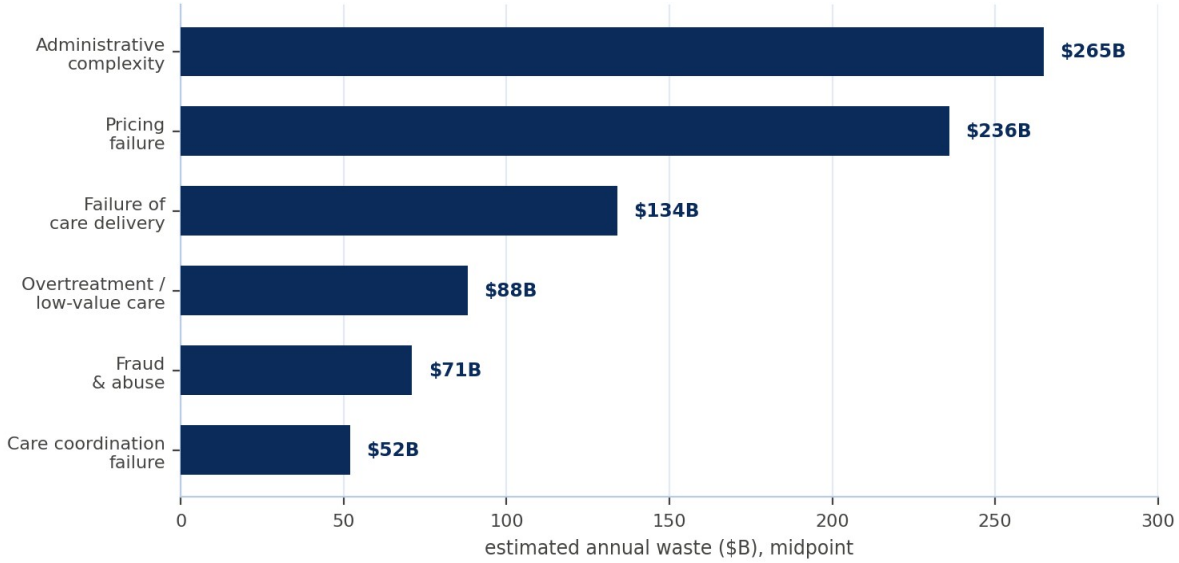

Figure 3. Estimated annual waste in U.S. healthcare, by category (midpoint of published ranges). Administrative complexity and pricing failure dominate; failure of care delivery and care-coordination failure are largely operational. Source: Shrank et al., JAMA 2019.

Administrative complexity alone accounts for roughly $265 billion. Failure of care delivery and failure of care coordination, the operational categories most directly tied to how the work is run, together add well over $150 billion, and the authors note these are where intervention is most feasible. This reframes the investment thesis. The central problem is not scarcity; the money, the doctors, even the technology already exist, and they are coordinated badly. Such waste does not shrink when you spend more. It shrinks when you operate better, which is an engineering and logistics problem, and a market.

II. Where it breaks: operational failure at the point of care

Macroeconomic waste is an abstraction until it is traced to the floor of a hospital, where it appears as four concrete, measurable failures: a collapsing workforce, a documentation burden that drives clinicians away from patients, a capacity-and-flow system that jams, and a steady stream of preventable harm. Each is quantifiable, each is expensive, and each is operational.

II.1 The workforce collapse

The nursing workforce, the largest and most load-bearing part of the hospital labor force, is turning over at an unsustainable rate. Registered-nurse turnover reached 17.6% in 2025, up from 16.4% the year before, at an average replacement cost of $60,090 per nurse; the average hospital loses $4.2-6.2 million a year to RN turnover alone, and every one-percentage-point change is worth roughly $262,000-295,000 (NSI Nursing Solutions, 2026). In the hardest-hit units, telemetry, step-down, and emergency, five-year cumulative turnover exceeds 110%, meaning those departments **replace their entire nursing staff in under five years**. Roughly 800,000 RNs reported an intent to leave the workforce by 2027 (NCSBN, 2022).

Consider why this happens. Nursing is the largest profession in healthcare, roughly 4.5 million registered nurses who deliver the great majority of direct patient care, and yet, in the hospital's own accounting, nursing is not revenue. Its cost is bundled into the per-diem room-and-board rate rather than billed as a service, which renders the value of nursing invisible on the ledger and books the nursing unit as a cost center. Because nursing is at once the single largest labor cost in the building, roughly a quarter to nearly half of the operating budget, and the line that looks like pure expense, it becomes the most attractive thing to optimize, which in practice means to cut. **The largest, most load-bearing workforce in American healthcare is the one no one thought to ask what it needs.**

This is the false economy at the heart of hospital staffing. Cutting nursing positions looks like immediate savings, but a chronically short unit cannot function, so the gap is backfilled with agency and travel labor at roughly $91 an hour against $59 for an employed nurse (two to three times staff cost during shortages), paid to a clinician with no lasting connection to the organization or its mission (NSI, 2026). Understaffing does not save money; it converts a fixed, loyal workforce into a permanent, premium-priced, rotating one. The vacancy data make the chronic shortage plain. RN vacancy peaked near 17% in 2022 and remains in double digits at roughly half of U.S. hospitals, with high-acuity units higher still. Many units are simply, **always short**. Labor already accounts for about 56% of hospital operating cost, and the bill for running this way grew more than $42 billion between 2021 and 2023 (AHA). The crisis is not cyclical. It is driven by the conditions of the work, the chaos, the cognitive overload, the time stolen from patients by administrative burden, and it will not be solved by recruiting harder into a leaking bucket.

Understaffing is also unsafe, in ways the evidence has settled. The work of Aiken, Lasater, and colleagues shows that each additional patient added to a nurse's average workload raises the odds of 30-day patient mortality by roughly 16%, and that lower staffing is directly associated with more patient falls, more hospital-acquired pressure injuries, more failure-to-rescue, and more missed care. The mechanism is not mysterious: **patients fall because there is no one free to help them to the bathroom, and pressure injuries form because there is no one free to turn them and get them out of bed**. These are not lapses of judgment to be fixed with another performance-improvement project; they are capacity failures. The reflexive institutional response, wait for falls to rise and then launch a new initiative to fix a problem only more hands can solve, treats the symptom while protecting the cause.

The cost of running this way is not borne only at the bedside. The nurse manager responsible for keeping the unit staffed spends upwards of 50% of their time building and reworking the schedule:

soliciting requests, assembling the roster, publishing it, then re-managing it every single day as call-outs, census swings, and acuity changes blow holes in it. A roster that takes tens of hours to build across two weeks must then be defended shift by shift, by hand. That is half of a senior clinical leader's capacity spent on a combinatorial math problem, the kind of constraint-satisfaction task software solves in seconds, and every hour of it is an hour not spent coaching, retaining, and protecting the team, which is itself a documented driver of turnover (NurseGrid; nurse-leader span-of-control literature, 2024). **The people fighting hardest to keep units safe are buried in the manual labor of doing so.**

The crisis is not confined to nursing. Physician burnout, though down from its 2021 peak, still affected 41.9% of doctors in 2025 (from 48.2% in 2023), with primary care and emergency medicine at the high end (AMA/Stanford Medicine, 2025). Replacing a single departing physician costs an estimated **$500,000 to $1,000,000** in recruitment, locum coverage, and lost billings, often two to three times the physician's salary, and burnout is estimated to cost the U.S. health system about $4.6 billion a year, predominantly through turnover and reduced clinical hours (Annals of Internal Medicine, 2019). Across both professions the pattern is identical: the scarcest, most expensive labor in the system is burned out by the conditions of the work and then expensively replaced, a self-inflicted operating cost a better-run hospital would not incur.

II.2 The documentation trap

A defining study of clinician time found that for every hour physicians spend with patients, they spend roughly two additional hours on the EHR and desk work, plus one to two hours of after-hours documentation, so-called pajama time (Sinsky et al., Annals of Internal Medicine, 2016). The bedside nurse's version is worse, because the nurse has become the **integration layer between every system in the building**, reconciling the EHR, the medication record, the lab system, the staffing board, the supply room, and the phone, performing in their head the data joins a database should perform in milliseconds. Every minute spent reconciling systems is a minute not spent on care, and it is the single most cited driver of burnout.

This is the operational tragedy in miniature: the most expensive, most highly trained, and scarcest resource in the hospital spends the majority of its time on analog coordination work that software should have absorbed. It is also the clearest opportunity. Documentation and coordination are the kind of structured, rule-bound, repetitive work that automation does well, and the kind that, once removed, returns a clinician to the bedside without hiring a single additional person.

II.3 Capacity and flow failure

Hospitals are at once too full and badly flowed. National inpatient occupancy rose from about 64% before the pandemic to 75% by 2024, not because demand surged but because the supply of staffed beds **fell by roughly 128,000 (about 16%)**, the equivalent of closing 160 average hospitals, as the workforce shrank (UCLA/JAMA Network Open, 2025). That same analysis projects national occupancy could reach 85% by 2032, a threshold, as Section IV explains, at which delay becomes self-reinforcing.

The result is visible in the emergency department, where boarding admitted patients in hallways and corridors has become, in the words of the American College of Emergency Physicians, a public health crisis. The share of ED encounters boarded more than three hours rose from 22%

in 2017 to 36% in 2024 (EDBA), and each additional hour of boarding measurably raises the risk of patient deterioration. Upstream, roughly a quarter of all inpatient days are avoidable (patients medically ready to leave who cannot, mostly because post-acute placement and internal processes fail), amounting to about 10.8 million avoidable inpatient days nationally each year (Advisory Board). Average length of stay rose 19% between 2019 and 2022 (AHA). None of this is a clinical-quality problem. It is a logistics problem: the right patient is not in the right bed at the right time, and no one in the building has a real-time model of the whole.

II.4 Preventable harm

When a system is overloaded and uncoordinated, it harms people. The most rigorous recent estimate found that **23.6% of hospital admissions — nearly one in four — involved at least one adverse event**, and that 22.7% of those events were preventable (Bates et al., NEJM 2023). A federal review reached a similar figure: a quarter of Medicare patients experienced harm, 43% of it preventable (OIG, 2022).

Those harms carry two prices, what the hospital pays and what the patient pays, and since 2008 Medicare has classified the worst of them as never events and hospital-acquired conditions it will not reimburse, so the hospital absorbs the full financial cost of harm it caused.

Never event	U.S. incidence / year	Cost to the hospital per event	Toll on the patient
Inpatient falls (with injury)	700,000-1,000,000 falls; 30-35% injure	≈$62,000; adds 6+ inpatient days	Fractures, head trauma, death; CMS pays nothing
Pressure injuries (HAPI)	≈2.5 million cases	≈$21,800 avg; $75,000-$150,000 if severe	Up to ~60,000 deaths/yr; sepsis; 1.5-2× readmission
CLABSI (central-line infection)	Tens of thousands	≈$48,000	≈12-25% attributable mortality
CAUTI (catheter UTI)	Hundreds of thousands	≈$13,800	Sepsis; prolonged stay
All hospital-acquired infections	≈687,000 infections	**$28-45B nationally**	**≈72,000 deaths/yr (CDC)**

Sources: AHRQ and Joint Commission (falls); Padula et al., International Wound Journal 2019/2021 (pressure injuries, ≈$26.8B nationally); CDC and AHRQ (infections); CMS Hospital-Acquired Condition no-pay policy (2008-).

Read together, the national bill is on the order of $20 billion a year for inpatient falls, $26.8 billion for pressure injuries, and $28-45 billion for infections, and the human ledger runs to tens of thousands of deaths and millions of injuries a year. The decisive point ties this section back to the workforce: **these are not random misfortunes but the measured downstream of how the operation is staffed and run**. They are the falls and pressure injuries Section II.1 traced to too few hands, plus the infections that climb when monitoring and protocol enforcement slip under load. So the hospital pays twice, first to understaff and then for the harm understaffing produces, while Medicare's no-pay rules and readmissions penalties convert every preventable event directly into lost revenue. These are the events that real-time monitoring, protocol enforcement, and closed-loop workflow are built to prevent.

One caution on the human numbers: the widely repeated claim that medical error is the third leading cause of death (250,000-440,000 deaths a year) is methodologically contested and should be treated with care. The defensible point is not a precise national body count but the

convergent, well-measured findings above: tens of thousands of deaths from specific, trackable, largely preventable events, on top of the one-in-four admissions that carry an adverse event.

III. Why the technology hasn't worked

If the problem is operational and the money is enormous, why hasn't two decades of healthcare technology fixed it? Because the technology solved the wrong problem. It digitized records and optimized fragments while the operation as a whole stayed uncoordinated. Understanding how the existing tools fall short is the key to seeing the opening.

III.1 The $36 billion EHR experiment

The 2009 HITECH Act spent more than $35 billion in federal incentives to drive electronic health record adoption from about 3% of hospitals in 2008 to 96% of non-federal acute-care hospitals by 2017 (Health Affairs, 2016). It was, on its own terms, a spectacular success, and it did not bend the cost curve, improve outcomes, or reduce burnout. It made them worse. The reason is foundational: the EHR was designed as a **billing and legal record**, a digital filing cabinet rather than an instrument for running the operation. It captures what happened for the purpose of getting paid and defending against liability. It does not decide what should happen next, bring the right task to the right person, or hold a live model of the unit. Digitizing an analog filing cabinet yields a faster filing cabinet, not a coordinated hospital.

Worse, the EHR pushed new work down onto clinicians, who became the data-entry staff for a system optimized around reimbursement. The productivity paradox of health IT, enormous investment and near-universal adoption with no measurable system-level return, is the single most important and most expensive lesson in the sector, and direct evidence that more record-keeping software is not the answer.

III.2 Point-solution sprawl

Into the gap the EHR left rushed a fragmented market of point solutions, each solving one slice. Capacity and command-center vendors (GE HealthCare, TeleTracking, LeanTaaS, Qventus, Hospital IQ) watch beds and transfers. Ambient documentation vendors (Microsoft/Nuance DAX, Abridge, Nabla, Suki) transcribe the visit. Scheduling and staffing platforms (ShiftMed, IntelyCare) optimize labor cost. Communication apps (Symplr/Halo, Voalte, Spok) move messages. Each can show a real, narrow result, yet each is a local optimum that leaves the system as a whole uncoordinated, because none of them holds the cross-system model the others would need to plug into.

This is the deeper failure mode, and it is a law of operations: **optimizing every component in isolation does not optimize the system**. A faster OR schedule that feeds a jammed ward, or an ambient scribe that documents a visit the bed-management system never sees, moves the bottleneck without removing it. The hospital does not need ten more gardens to tend; it needs one layer that consumes the gardens and coordinates across them.

III.3 The black-box AI trust collapse

The current wave of healthcare AI arrives at a moment of collapsing clinician trust, and for good reason. The most-deployed predictive model in American hospitals, the Epic Sepsis Model, was externally validated at an area-under-curve of just 0.63 against its advertised 0.76-0.83, missed roughly two-thirds of sepsis cases, and generated enough false alerts to drive serious alarm fatigue. It became the canonical example of why narrow, opaque, vendor-validated AI fails in clinical practice. As vendors pour more black-box models into the market, clinicians trust them less, and regulators have noticed.

This trust collapse is itself a market signal. The 2026 regulatory environment (the Joint Commission's Responsible Use of AI in Healthcare certification and the ONC's HTI-1 algorithm-transparency rules) is moving decisively toward requiring that algorithmic decisions be explainable and auditable. A deterministic, citation-backed, transparent system is not just safer; in 2026 it is becoming a procurement requirement. The market is rotating away from the black-box approach most AI vendors are built on.

III.4 The category error

Beneath every failure above is one conceptual mistake. The industry has spent two decades trying to optimize the productivity of individual human beings, squeezing more charting, more patients, more throughput out of each clinician, when the science of operations says the lever is **system flow, not individual effort**. You cannot fix a queuing system by making each server work harder; past a point, that makes it worse. Hospitals have been running the building's most expensive, least scalable resource, trained humans, at ever-higher utilization, and calling the resulting burnout a personal-resilience problem. It is not. It is a design problem, and Section IV explains why it was always going to fail.

III.5 The misplaced AI bet and the credibility gap

There is a final reason the technology has not worked: it has been aimed at the wrong target. The current wave of healthcare AI is concentrated on the physician's diagnostic and decision-making work, the imaging reads, clinical decision support, and ambient scribing, a domain already among the most optimized in the hospital and where, for most tasks, the experienced human still outperforms the model. The settled use cases in which AI reliably beats a clinician are few and narrow. Meanwhile, hundreds of computational and logistical processes are executed by hand every day across the building, including bed assignment, staffing math, throughput forecasting, registry assembly, compliance evidence, and supply reconciliation, and these have barely been automated at all. The investment went where the human still wins and skipped the work a machine would do better.

Compounding the misallocation is a credibility problem the sector underrates. Clinician attitudes toward AI are not uniformly positive, and the skepticism is concentrated where it matters: a nationwide survey of 991 physicians and 1,714 nurses found that perceived risk significantly reduced nurses' intention to use AI, with reliability, transparency, and liability for error the dominant concerns (JMIR, 2025). The Epic Sepsis Model, validated externally at an AUC of 0.63 and missing roughly two-thirds of cases, is the cautionary tale every clinician now knows. And the incoming workforce is the most skeptical of all: Gen Z's enthusiasm for AI is falling while anger

rises (31% say AI makes them angry, up nine points in a year; only 22% are excited, down from 36%), and fewer than three in ten Gen Z workers trust AI-assisted work (Gallup, 2025-2026). **A platform that depends on clinicians trusting a black box is fighting the tide.**

This is why an operational-control approach is also lower-risk. It does not ask anyone to trust a novel model; it replicates **time-tested logistical analysis** (queuing theory, Erlang staffing math, constraint solving) proven for decades in aviation, the military, and manufacturing, and delivers it on technology surfaces hospitals have already adopted. Reproducing a settled method on a familiar surface is its own proof of concept: the mathematics is not in question, only its application. The output is deterministic, explainable, and auditable, the property clinicians say they require and the one regulators are now mandating. The opportunity is not smarter AI for the doctor; it is basic, proven automation for the operation.

IV. The idea that isn't new: operational and logistical control

The solution healthcare needs is not novel. It is the discipline every other complex, high-consequence system (military, aviation, logistics, manufacturing, telecommunications) adopted decades ago and that healthcare, almost uniquely, never did: centralized operational and logistical control, grounded in the mathematics of flow. The ideas are a century old. Applying them to the hospital is the opportunity.

IV.1 What the military solved

Operations research was born in World War II, when the Allied militaries put civilian scientists to work on the logistics of war: how to route convoys, allocate scarce equipment, plan maintenance, and coordinate enormous, variable, life-or-death systems under uncertainty (Britannica; INFORMS). The methods worked so well that the postwar military institutionalized them, and they spread into industry as the foundation of modern supply chains, airline operations, and just-in-time manufacturing. The military's genuine, hard-won expertise is not weaponry; it is logistics: moving the right resource to the right place at the right time, at scale, reliably, when the cost of failure is measured in lives. That is the capability a hospital lacks.

IV.2 The 85% wall

The core mathematics of operational control is queuing theory, and it explains the hospital's failure with unusual precision. In any system where work arrives with variability and is served with variability, whether patients to beds, to nurses, or to operating rooms, waiting time is a nonlinear function of utilization. As utilization rises toward 100%, delay does not increase proportionally; it increases **exponentially**. A foundational hospital study (Bagust et al., BMJ 1999) found that bed shortages and crisis days become effectively inevitable above roughly 85% occupancy, and the principle is general to any high-variability queue.

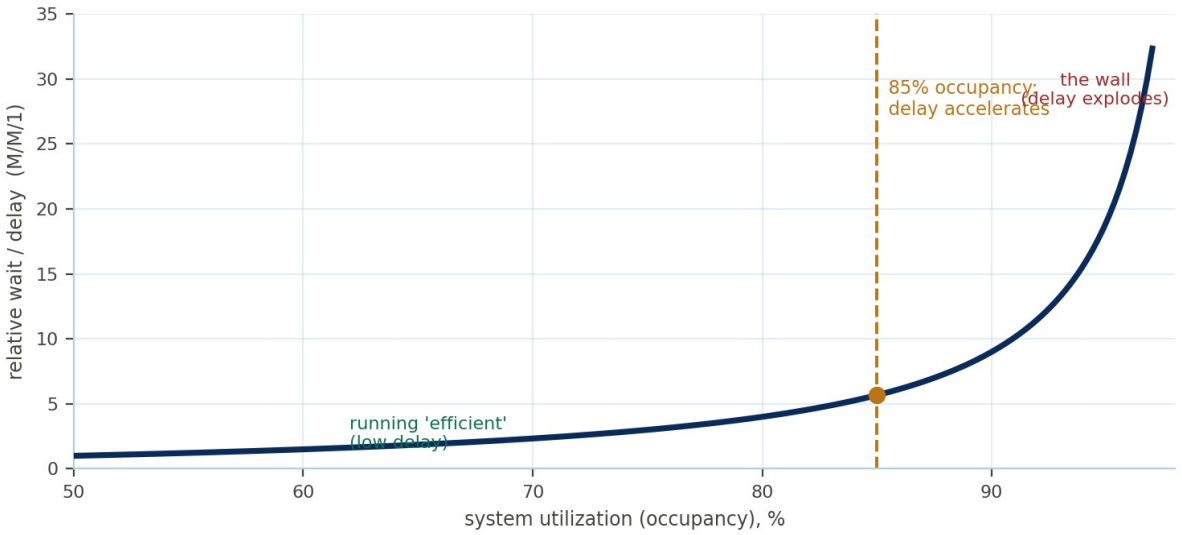

Figure 4. The mathematics of running hot. In a variable queue, average delay rises gently with utilization, then explodes as the system approaches capacity. Hospitals chase efficiency by pushing toward the right of this curve; the curve is why "it works until it doesn't." Illustrative M/M/1 relation; the 85% threshold is the classic hospital finding (Bagust et al., BMJ 1999).

This single curve indicts the dominant management philosophy of American hospitals. Administrators, taught to treat empty beds and idle staff as waste, push utilization ever higher in the name of efficiency, and they operate, by design, on the steep part of the curve, where a single admission surge or a few call-outs tips the unit into gridlock, diversion, and harm. The slack that queuing theory says a variable system requires to function is treated as the first thing to cut. Running hot looks efficient on a spreadsheet and is catastrophic in practice, because the spreadsheet is linear and the system is not. This is the rigorous, non-rhetorical form of the claim that the current approach works until it doesn't: a theorem, not an opinion.

IV.3 Automate the analog, unlock the human

If you cannot fix a queuing system by working the humans harder, what do you do? You manage the system for flow, and you remove from the humans every task that does not require a human. The hospital is full of analog coordination work, reconciling systems, hunting for data, updating boards by hand, assembling registry submissions, remembering what this week's protocol requires, all of which consumes the clinical workforce and all of which software can absorb. **Automating the analog frees the system's most precious resource, human beings**. A nurse relieved of the integration-layer burden is, in effect, additional capacity created without additional hiring, the only kind a workforce-constrained system can actually add.

This inverts the standard model. The point is not to make documentation faster so the clinician can see more patients per hour; that is the productivity trap again. It is to delete the analog work so the system runs on flow and the human is returned to judgment, presence, and care, the things only a human can do and the things the current system has been quietly rationing. Automation here is not labor replacement but labor liberation, and it is the lever that attacks burnout, turnover, capacity, and safety at once, because all four share the same root cause.

This is where the abstract argument becomes a number. The platform's design target is to remove on the order of **40% of the analog workload from the bedside nurse and half of the**

coordination overhead from the charge nurse and manager, the documentation, the data-shuffling, the manual board updates, the hunting across systems. A nurse who spends 40% less time fighting the system, and a manager freed from the scheduling grind that, as Section II.1 showed, consumes roughly half of their time, together amount to a near-doubling of usable clinical capacity from the same roster, capacity created without hiring a single additional full-time employee, the only kind a short-staffed, agency-dependent hospital can realistically add. Reductions of this magnitude are documented in the literature: ambient AI cut clinician burnout by 40% at Mass General Brigham and documentation latency by 81% at Cedars-Sinai, nurses lose roughly a fifth of every 12-hour shift to documentation and reach the bedside for only about a third of their time, and the company's own pilot observed a 35% documentation-time cut. The point is emphatically not to replace the nurse. Problems in hospitals can only be solved by people; what the platform does is give those people back the time and resources to do the job they came to do and currently cannot reach.

IV.4 The Navy model

There is a cultural half to this that the data alone does not capture, worth stating plainly because it is the difference between the operational-control thesis and ordinary efficiency software. The cleanest model is the Navy. If the Navy expects an individual to perform at an elite level, it builds the entire system around them, the training, the tools, the precision instruments, the doctrine, so that elite performance becomes the reliable default. The investment and the expectation move together. The High Reliability Organization framework that healthcare aspires to was borrowed directly from naval aviation and nuclear power; healthcare adopted the aspiration and, by its own outcome data, not the practice.

Hospitals invert the Navy's bargain: they expect new clinicians to perform flawlessly and then place the entire burden of getting there on the individual, treating burnout as an acceptable cost so long as new hires keep arriving. An operational-control platform is the tooling half of the bargain the Navy honors and healthcare breaks, the precision instruments that let an ordinary clinician perform reliably without first becoming a systems engineer. It does not replace investment in people; it makes that investment pay off instead of leaking out through friction and turnover. That is the philosophy underneath the software, and it is why the right architecture is an operating system for the hospital, not another app.

V. Market positioning and opportunity

Sections I-IV are an objective reading of public data. This section is interpretive: it lays out where the market is, why the white space exists, how large it is, and how Nightingale OS is positioned to occupy it.

V.1 The white space: a system of action

Hospital software today is a stack of systems of record, destinations a clinician navigates to enter data. The unoccupied category is the system of action that sits on top of them: a governed operating layer that consumes the existing silos, holds a live model of the operation, predicts the predictable, and surfaces the next action, already populated, to the accountable person, on the right device, at the right moment. It is the difference between a filing cabinet and an operating

system. No incumbent ships this, because building it requires starting from a cross-system operational model and a governance substrate rather than from a record, a different architecture and a different business model than any incumbent has.

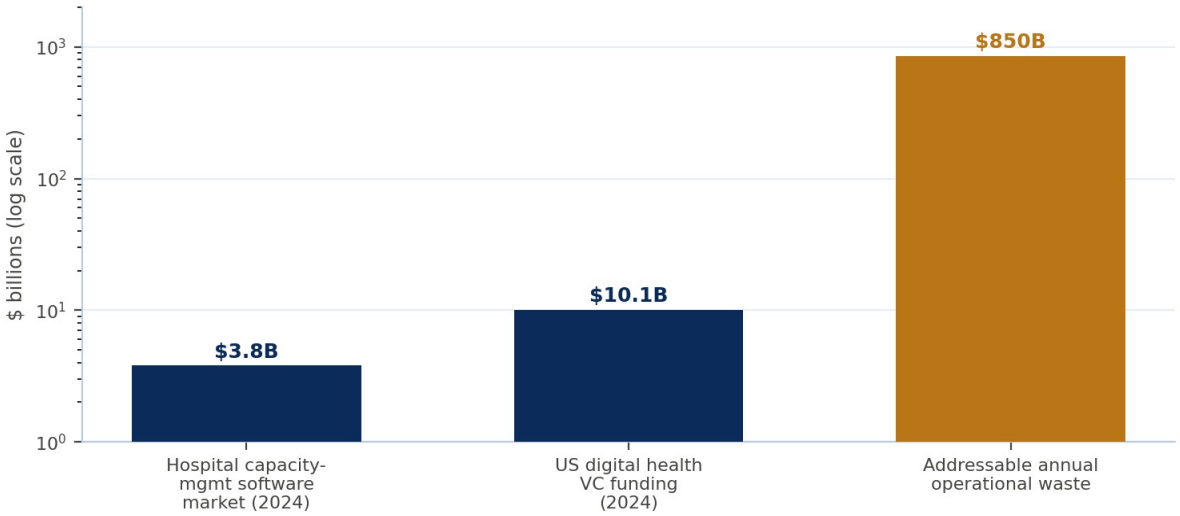

Figure 5. The market gap. Current operations tooling and even total digital-health venture funding are small against the operational waste they are meant to address. Sources: Grand View Research (capacity-management market, 2024); Rock Health (2024 funding); Shrank et al., JAMA 2019 (waste). Log scale.

V.2 The integration thesis: an agnostic adapter

The shape of the solution follows from the diagnosis. If the problem is that a hospital's dozens of systems do not talk to each other and the staff are the integration layer, then the product is an agnostic adapter that lets the applications already in use work together: a governed layer that reads from the EHR, the lab and pharmacy systems, the bed board, the staffing system, RTLS, and telemetry through open standards (FHIR R4, HL7 v2, DICOM, and vendor APIs across roughly two dozen integration categories) and coordinates across them. This is, in effect, the answer to the long-term IT-networking and interoperability problem hospitals otherwise face. It is not another walled garden to feed but the connective layer that makes the gardens a hospital already owns finally useful together.

The contrast with the incumbent model is stark. The EHR systems running American hospitals are old, architecturally little more sophisticated than networked spreadsheets, and they are **effectively impossible to update** without building an entirely new product and then persuading the hospital to buy it. An Epic-class installation runs $10-30 million, and $100-300 million or more for large systems, takes years, and carries 15-22% of license cost in annual maintenance plus a six-figure internal IT team (industry estimates, 2025-2026). A coordination layer built as a cloud-based web application inverts all of that: it deploys in weeks rather than years, it improves continuously without a forklift upgrade, and the only hardware a hospital must purchase and install is the smartboards. That is a different cost structure, a different risk profile, and a different speed than anything in the system-of-record world.

V.3 Charlie: the experience layer and the demand-side moat

Above the coordination layer sits the experience layer, the part clinicians and patients actually touch, where the platform turns from infrastructure into preference. In the product it is called Charlie, and it has two faces. The first finally fixes the bedside information problem. Most American hospitals still run their units on dry-erase whiteboards updated by hand, and patients are routinely left in the dark about their own care: who their nurse is, what the day's goals are, what comes next. A governed smartboard, clinical when staff are present, patient-private when idle, a family wayfinding surface besides, replaces the manual whiteboard and the taped-up fliers with a live, self-updating nervous system for the room. The smartboard is the only hardware a hospital buys, it is revolutionary on its own, and the technology has existed in other industries for years; it was simply firewalled out of healthcare.

For clinicians, Charlie Assist surfaces the next task in order of clinical urgency, documents ambient conversation into structured notes, and shows the care team, the day's goals, and overdue meds and reassessments, all behind a $k \geq 5$ privacy floor that never surveils the individual. The pilot modeled it returning roughly 2,457 nursing hours a year and cutting documentation time 35%. The strategic point is not the hours; it is what working this way does to retention. **Once a clinician has worked a shift with Charlie, the old way becomes unthinkable**. A clinician who will not go back is one who chooses, and stays at, a hospital that runs it. In a labor market where replacing one nurse costs about $60,000 and one physician $500,000 to $1,000,000, that retention effect is no soft benefit. It is a direct, compounding attack on the largest controllable cost in Part II, and a recruiting advantage a competitor on the old stack cannot match.

That advantage is deepened by something competitors will not build: a **trauma-informed architecture**. The platform is designed against the SAMHSA trauma-informed principles, treating the psychological safety of the workforce as an engineering requirement rather than a poster in the break room: surfacing help without surveillance, protecting worker voice, and refusing to optimize a metric at a person's expense. For investors, this is an investment in workforce readiness, a system that makes a new nurse reliable and a veteran nurse willing to stay. A labor-optimization vendor, whose business is selling the buyer lower labor cost, cannot credibly ship worker protections that constrain its own customer, so a humane, trauma-informed operating layer is a differentiator defensible precisely because the incumbents' business model forbids it.

The second face, Charlie Companion, points at the patient and reaches the one lever hospital software almost never touches: demand. Almost automatically it delivers something close to concierge service: the care team's names and photos, the day's goals in plain language, comfort and pain input, education matched to the diagnosis, and a channel for family to reach the team. Patients who can see their plan, know their nurse, and feel informed feel safer, and in an era of public hospital ratings and real consumer choice, experience is a reason patients choose one system over another. **A platform patients prefer and clinicians refuse to give up pulls from both the demand side and the supply side at once**. It is a combination almost unheard of in enterprise hospital software, the kind of differentiator that compounds: the hospitals that adopt it become the ones patients seek out and clinicians want to work for.

V.4 What hospitals buy, what it costs, and what it misses

The landscape is fragmented by design: every category solves a slice and none provides the unified, governed, real-time layer. That fragmentation is the opportunity and, given the integration and trust required to span it, the barrier to entry.

Category	Representative players	Structural limitation
EHR / system of record	Epic (42% of hospitals, 55% of beds), Oracle Health/Cerner, MEDITECH	A billing and legal record, not an orchestration layer; monetizes breadth of the record, not cross-system flow.
Capacity / command center	GE HealthCare, TeleTracking, LeanTaaS, Qventus, Hospital IQ	Narrow domain (beds, transfers, OR); dashboards over a few feeds; no governance substrate; cloud-dependent.
Ambient documentation	Microsoft/Nuance DAX, Abridge, Nabla, Suki	Cloud-first by construction; cannot serve VA, federal, rural, or behavioral-health buyers; transcribes the visit, does not run the operation.
Staffing / workforce	ShiftMed, IntelyCare, scheduling vendors	Optimize labor cost for administration; structurally opposed to shipping worker protections that constrain their own buyer.
Quality / safety	HAC and infection-tracking dashboards	Retrospective reporting; no predictive, closed-loop, or governance layer.
System of action (white space)	**— unoccupied —**	**The vendor-agnostic, real-time, governed layer over everything. The gap and the moat.**

Hospitals already spend heavily across that fragmented stack, and the prices are not small.

Category	Typical implementation + annual cost	What it does — and misses
EHR / system of record	$10M-$300M+ to install; ~15-22% of license/yr maintenance plus a $150K-$500K IT team; ~$1,000-$2,500 per bed/yr	Holds the record. Old and monolithic — cannot be meaningfully updated without buying a new product.
Capacity / command center	$150K-$500K+ per facility/yr	Watches one domain (beds, OR, transfers). Another dashboard, another silo.
Ambient documentation	$100-$300 per provider/month	Transcribes the physician visit. Cloud-only; does not touch the operation.
Staffing / scheduling	$50K-$200K per facility/yr	Optimizes labor cost. No flow model, no governance.
Secure messaging	$20K-$80K per facility/yr	Moves messages between silos. No coordination.
Quality / safety dashboards	$30K-$150K per facility/yr	Reports harm after the fact. No closed loop.

Pricing from industry implementation estimates and Nightingale market research; figures are directional and negotiated per facility.

The deeper cost is what this stack does to the people inside it. Each product is a separate contract, a separate login, and a separate silo, and the staff become the integration layer, shuttling data from one program to another by hand and re-keying the same facts into the EHR, the command center, the scribe, the scheduler, and the messaging app. The industry's push toward single-solution products compounds the problem instead of solving it: every new best-of-breed

tool is one more island. And because hospital budgets are finite, a one-off product that moves the needle only a little earns limited adoption and quiet abandonment. This is why systems accumulate an ever-taller stack of advanced solutions that work sometimes, while the thing that would make them work together is never bought, because no one sells it.

This is the layer the industry skipped. **The opportunity is not a better point solution; it is the missing coordination layer** that lets the applications a hospital already owns operate as one system. That is a category, not a feature, and it is empty.

V.5 Modular by design: bundles and the Epic-shop wedge

Because the platform is a coordination layer rather than a monolith, it is modular by design: a catalog of independently activatable bundles organized into product lines, any of which a hospital can switch on without ripping out anything it already runs. This dissolves, at once, the two objections that kill most hospital software deals: rip-and-replace risk and the budget ceiling. A hospital need not choose between Nightingale OS and Epic; it can adopt the single bundle whose return is cleanest, prove it on its own data, and expand from there.

The compliance-and-audit bundle (internally, Owl) is the sharpest illustration. It continuously assembles, from data already flowing through the platform, the accreditation and ratio-law evidence a surveyor or arbitrator would demand: a tamper-evident, real-time compliance record. **No incumbent does this today**; compliance remains a retrospective, manual scramble staged the week before a survey. A bundle that turns it into a continuous, provable artifact is valuable even to a hospital fully committed to Epic, and rich with adjacent possibility: automated registry submission, regulatory change-watch, penalty-exposure tracking, each a product in its own right. Modularity is how a coordination layer lands inside an incumbent shop, not by replacing the record but by doing the things the record was never built to do.

V.6 Market sizing and unit economics

The addressable market can be sized two ways. Top-down, the operational waste the category attacks is on the order of $850 billion a year, while the tools aimed at it remain small: a roughly $3.8 billion hospital capacity-management software market (Grand View, 2024) inside a digital-health sector that drew $10.1 billion of U.S. venture funding in 2024 (Rock Health). Bottom-up, there are about 6,100 U.S. hospitals and 913,000 staffed beds (AHA, 2025); a per-bed operations-and-governance platform represents a **$500 million to $2.5 billion annual revenue opportunity** in U.S. acute care alone, extending toward $3-5 billion with adjacent settings and accreditation use cases.

Layer	Estimate	Basis
TAM	$0.5B-$2.5B ARR (to $3B-$5B with adjacencies)	~913,000 U.S. staffed beds at a per-bed operations + governance price; AHA 2025.
SAM	$200M-$800M ARR	~2,500-3,000 acute-care hospitals where the fit is acute: ratio-law states, non-Epic, federal/VA/rural excluded from cloud vendors.
SOM (Year 2)	$1.5M-$8M ARR	10-20 community hospitals in beachhead states at ~$150K-$400K ARR each.

The unit economics are anchored by the company's digital pilot, which modeled a 274-bed hospital and projected roughly $2.0M in net annual value against a $59,000 platform cost, about $7,500 of modeled value per bed at a 34.9× return (modeled, not observed). Read against the cost of the status quo, the ROI case is concrete: the average hospital loses $4.2-6.2M a year to nurse turnover, one hour of ED boarding reduction can recapture five figures per day, and Medicare's penalty programs convert preventable harm directly into lost payment. The platform's job is to surface those scattered, invisible losses as a single managed, financeable opportunity.

V.7 The closure crisis and the social-impact market

The unsustainability documented in Part I is not distributed evenly; it is concentrated where it does the most human damage. Nearly half of rural hospitals, 46%, operate with negative margins, and 432 are currently vulnerable to closure; 18 closed or stopped offering inpatient care in the past year, and 182 have done so since 2010. Beyond outright closure, communities are losing core services: 293 rural hospitals dropped obstetrics between 2011 and 2023, and 424 ended chemotherapy (Chartis, 2025). When a rural hospital closes, the nearest emergency care can move an hour away, and the economic anchor of the town goes with it.

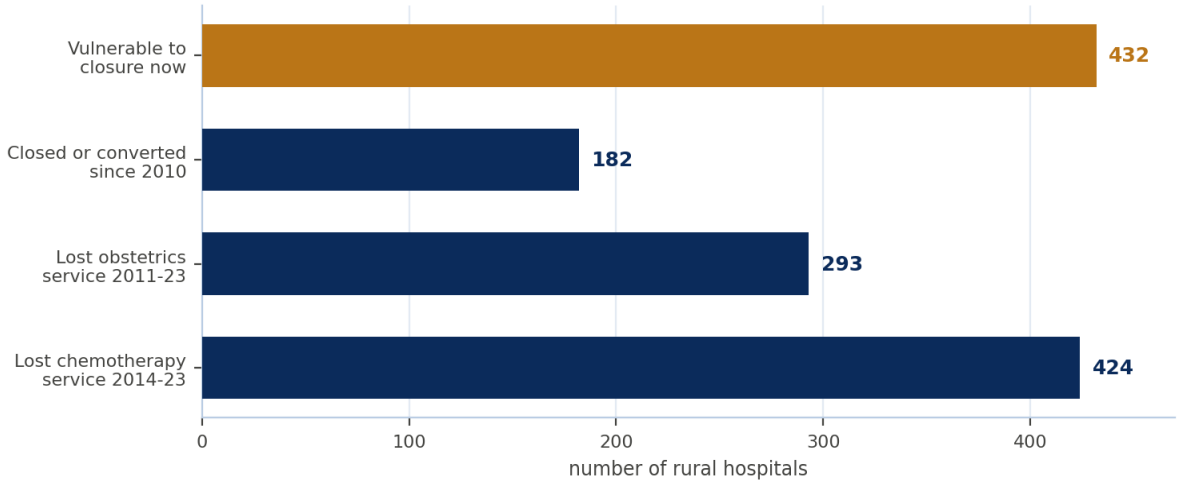

Figure 6. The rural safety net under pressure. 46% of rural hospitals run negative operating margins. Source: Chartis, 2025 Rural Health State of the State.

This is, bluntly, where an operational-control platform matters most, and it is also a defensible market segment. A facility losing money cannot cut its way to survival on a skeleton staff; what it needs is to **eliminate waste and recover margin without adding headcount**, which is what coordinating the operation does. A digital twin that knows the standard and adapts it to the resources a hospital actually has lets a 25-bed critical-access facility run to a level of reliability it could never staff a consultant or a registry team to reach. Critical-access and rural hospitals are also structurally excluded from cloud-only vendors by data-residency and connectivity constraints, so an on-premise, low-hardware, weeks-to-deploy layer reaches buyers the incumbents cannot. The result is a rare alignment: the segment with the most acute need, the segment the competition cannot serve, and a genuine instrument of social change, keeping care in the underserved places that can least afford to lose it.

V.8 Why now

Three forces converged in 2026 to make an operational-control platform both required and buildable. First, regulation turned transparency from a virtue into a procurement requirement: nurse-to-patient ratio laws are tightening with real financial penalties, the Joint Commission's Responsible Use of AI certification launched in June 2026, and the ONC's transparency rules now demand auditable algorithms, the very artifacts a deterministic, citation-backed system produces by design. Second, the collapse of clinician trust in black-box AI rotates the market toward transparent, governable systems. Third, the staffing crisis is now understood as structural rather than cyclical, so buyers want tools that retain staff by removing burden, not tools that squeeze them harder.

V.9 The Nightingale OS thesis

Nightingale OS is built to be the system of action: a governed operating layer that runs on-premise alongside any EHR, consumes the hospital's existing systems, holds a live operational model, and surfaces governed, cited, auditable action to clinicians, with worker protections and compliance proofs enforced in code rather than asserted in policy. It is the operational-and-logistical-control thesis of this analysis, instantiated: automate the analog coordination work, manage the system for flow, return the human to care, and produce the financial and regulatory receipts that make the value undeniable to the people who hold the budget. The objective case in Sections I-IV is that the industry needs this. The market case in Section V is that no incumbent is built to deliver it.

VI. Risks and honest disclosures

A serious analysis states its weaknesses. The following are the material risks to the thesis and to the company, presented plainly.

- **Execution and integration risk.** The hardest part of any cross-system layer is the integration. As of June 2026 the company's ADT interface is built and tested but no live hospital feed has flowed; the first design partner is the make-or-break validation.
- **Modeled, not measured.** The pilot's value figures are projections built on synthetic data and literature-based rates, not observed outcomes; clinical composites are labeled unvalidated pending advisor review.
- **Incumbent reflex.** “We already have Epic” is a real objection. The answer is that Epic is a record, not an orchestrator, and a worker-protective, vendor-agnostic layer is structurally off-brand for a system-of-record vendor — but distribution advantage is real.
- **Long, committee-driven sales cycles.** Hospital enterprise software averages ~12-month cycles with multiple decision-makers; adoption is slow even when ROI is clear.
- **Attribution risk on gain-share.** Healthcare outcomes have many causes; any value-based contract must define metrics, baselines, and measurement windows precisely.
- **Founder and stage risk.** A non-traditional solo founder at pre-seed; the domain-specific work is done, but the team and capital to scale must still be built.

VII. Conclusion

The United States runs a healthcare system that costs a fifth of its economy, delivers worse outcomes than nations spending half as much, and is on a fiscal countdown clock. The dominant share of the waste is operational: coordination, flow, staffing, documentation. Two decades and tens of billions of dollars of record-keeping software and point solutions have not touched it, because they digitized and optimized fragments instead of coordinating the whole. The discipline that fixes this class of problem is not new; it is the operational and logistical control the military perfected and that queuing theory makes rigorous. Apply it to the hospital, automate the analog work, manage the system for flow, and return the human to care, and the largest, most addressable inefficiency in the U.S. economy becomes a market. The system of action is the category. Building it well, and governing it by design, is the opportunity.

References

1. Centers for Medicare & Medicaid Services, Office of the Actuary. National Health Expenditure Data and Projections, 2023-2032 (2024). https://www.cms.gov/data-research/statistics-trends-and-reports/national-health-expenditure-data
2. Peterson-KFF Health System Tracker. How does health spending in the U.S. compare to other countries? (2025-2026). https://www.healthsystemtracker.org/chart-collection/health-spending-u-s-compare-countries/
3. The Commonwealth Fund. Mirror, Mirror 2024: A Portrait of the Failing U.S. Health System (Sept. 2024). https://www.commonwealthfund.org/publications/fund-reports/2024/sep/mirror-mirror-2024
4. OECD. Health at a Glance 2025 — United States. https://www.oecd.org/en/publications/health-at-a-glance-2025_15a55280-en.html
5. CDC/NCHS. Life Expectancy in the United States, 2024 (Data Brief No. 521, Dec. 2024). https://www.cdc.gov/nchs/data/databriefs/db521.pdf
6. CDC/NCHS. Maternal Mortality Rates in the United States, 2024 (Mar. 2026). https://www.cdc.gov/nchs/data/hestat/hestat113.htm
7. CDC/NCHS. Obesity and Severe Obesity Prevalence in Adults (Data Brief No. 508, 2024). https://www.cdc.gov/nchs/products/databriefs/db508.htm
8. Tanne JH et al. Avoidable Mortality Across US States and High-Income Countries. JAMA (2025). https://pubmed.ncbi.nlm.nih.gov/40126499/
9. Boards of Trustees, Federal Hospital Insurance and Federal Supplementary Medical Insurance Trust Funds. 2025 Annual Report (June 2025). https://www.ssa.gov/oact/trsum/
10. Shrank WH, Rogstad TL, Parekh N. Waste in the US Health Care System: Estimated Costs and Potential for Savings. JAMA. 2019;322(15):1501-1509. https://jamanetwork.com/journals/jama/fullarticle/2752664
11. Bates DW et al. The Safety of Inpatient Health Care. NEJM. 2023;388(2):142-153. https://www.nejm.org/doi/full/10.1056/NEJMsa2206117
12. Office of Inspector General, HHS. Adverse Events in Hospitals: A Quarter of Medicare Patients Experienced Harm (May 2022). https://oig.hhs.gov/reports/all/2022/
13. NSI Nursing Solutions. 2026 National Health Care Retention & RN Staffing Report. https://www.nsinursingsolutions.com/
14. National Council of State Boards of Nursing (NCSBN). Nursing workforce study (2022/2023). https://www.ncsbn.org/
15. NSI Nursing Solutions / Becker's Hospital Review. RN vacancy rate ~9.6% nationally in 2024, with roughly half of hospitals above 10% and a 2022 peak near 17%. https://www.beckershospitalreview.com/quality/nursing/hospital-nurse-turnover-vacancy-rates-by-year/

16. Aiken LH, Lasater KB, McHugh MD, et al. Nurse staffing and patient outcomes: each additional patient in a nurse's workload raises 30-day mortality odds ~16%; lower staffing is linked to falls, pressure injuries, failure-to-rescue, and missed care. JAMA / Medical Care / Lancet / JAMA Network Open, 2002-2025.
17. U.S. Bureau of Labor Statistics (OOH); ANA; NCSBN — registered nursing is the largest U.S. health profession (~4.5M RNs). Nursing cost is bundled into the per-diem room rate (a cost center, ~25-44% of operating budget): the 'invisible economics of nursing.' https://www.bls.gov/ooh/healthcare/registered-nurses.htm ; PMC5887158.
18. NurseGrid; Bradley University; HealthStream — nurse managers spend upwards of 50% of their time on scheduling (building, publishing, and daily re-working the roster); nurse-leader span-of-control integrative review, PMC10660758 (2024). https://nursegrid.com/blog/behind-the-schedule-the-overwhelming-ordeal-of-nurse-scheduling/
19. Sinsky C et al. Allocation of Physician Time in Ambulatory Practice. Annals of Internal Medicine. 2016;165(11):753-760. https://www.acpjournals.org/doi/10.7326/M16-0961
20. American Medical Association / Stanford Medicine. Physician burnout rate continues to decline, falling to 41.9% in 2025 (from 48.2% in 2023). https://www.ama-assn.org/practice-management/physician-health/physician-burnout-rate-continues-decline-falling-nearly-42 — and Han S et al., Estimating the Attributable Cost of Physician Burnout in the United States (~$4.6B/yr), Annals of Internal Medicine 2019.
21. American College of Emergency Physicians. Emergency Department Boarding & Crowding (statements, 2023-2024). https://www.acep.org/administration/crowding--boarding
22. Emergency Department Benchmarking Alliance. 2024 Data Performance Report. https://iepc.org/
23. Roy S et al. / UCLA. Hospital Occupancy and Staffed-Bed Supply. JAMA Network Open (2025). https://pmc.ncbi.nlm.nih.gov/articles/PMC11840646/
24. Advisory Board. Avoidable Hospital Days: Data Primer (CMS claims analysis). https://www.advisory.com/
25. CDC. National and State HAI Progress Report (2023). https://www.cdc.gov/healthcare-associated-infections/php/data/progress-report.html
26. AHRQ. Estimating the Additional Costs of Hospital-Acquired Conditions; Dykes PC et al., inpatient falls cost (2023). https://www.ahrq.gov/hai/pfp/index.html
27. Padula WV, Delarmente BA. The national cost of hospital-acquired pressure injuries in the United States (≈$26.8B). Int Wound J. 2019; and HAPI stage-wise mortality/readmission, Int Wound J. 2021 (up to ~60,000 deaths/yr). https://onlinelibrary.wiley.com/doi/abs/10.1111/iwj.13071
28. Centers for Medicare & Medicaid Services. Eliminating Serious, Preventable, and Costly Medical Errors — Never Events; Hospital-Acquired Condition no-pay policy (Deficit Reduction Act, 2008-). https://www.cms.gov/newsroom/fact-sheets/eliminating-serious-preventable-and-costly-medical-errors-never-events
29. CDC. Healthcare-Associated Infections: ~687,000 HAIs and ~72,000 associated deaths per year; CLABSI/CAUTI cost and attributable mortality. https://www.cdc.gov/healthcare-associated-infections/
30. Workload-reduction evidence: Mass General Brigham ambient-AI scribe pilot (~40% reduction in clinician burnout); Cedars-Sinai ambient nursing documentation (~81% reduction in documentation latency, 120+ nurses); systematic review of AI and EHR-related burnout, Frontiers in Public Health 2025. https://www.cedars-sinai.org/newsroom/artificial-intelligence-lightens-burden-on-nurses/
31. Kaufman Hall. National Hospital Flash Report (2024-2025); American Hospital Association, Cost of Caring (2025). https://www.kaufmanhall.com/insights-reports/national-hospital-flash-report
32. American Hospital Association. Fast Facts on U.S. Hospitals, 2025. https://www.aha.org/statistics/fast-facts-us-hospitals
33. Adler-Milstein J, Jha AK. HITECH Act Drove Large Gains in Hospital EHR Adoption. Health Affairs. 2017. https://www.healthaffairs.org/doi/10.1377/hlthaff.2016.1651
34. Rock Health. 2024 Year-End Digital Health Funding Report ($10.1B). https://rockhealth.com/insights/
35. Grand View Research. Hospital Capacity Management Solutions Market; Digital Health Market; AI in Healthcare Market; U.S. AI Medical Scribing Market (2024-2025). https://www.grandviewresearch.com/

36. Britannica; INFORMS. Operations Research: History (origins in World War II). https://www.britannica.com/topic/operations-research/History

37. Bagust A, Place M, Posnett JW. Dynamics of bed use in accommodating emergency admissions. BMJ. 1999;319:155-158. https://www.bmj.com/content/319/7203/155

38. Proudlove NC. The 85% bed occupancy fallacy: the use, misuse and insights of queuing theory. J Health Serv Res Policy (2020). https://journals.sagepub.com/doi/10.1177/0951484819870936

39. Makary MA, Daniel M. Medical error — the third leading cause of death in the US. BMJ. 2016;353:i2139 (contested). Critiques: Shojania & Dixon-Woods, BMJ Qual Saf 2017; Panagioti et al., BMJ 2019.

40. Attitudes, Perceptions, and Factors Influencing the Adoption of AI in Health Care Among Medical Staff. JMIR (2025). https://www.jmir.org/2025/1/e75343 — and Physicians' Attitudes Toward AI in Medicine, JMIR 2025;e74187.

41. Gallup. Gen Z's AI Adoption Steady, but Skepticism Climbs (2025-2026). https://news.gallup.com/poll/708224/gen-adoption-steady-skepticism-climbs.aspx

42. Chartis. 2025 Rural Health State of the State (432 rural hospitals vulnerable; 46% with negative margins; service losses). https://www.chartis.com/insights/2025-rural-health-state-state

43. Industry implementation-cost estimates for Epic and enterprise EHRs (2025-2026): ~$10M-$300M+ to install, 15-22% of license in annual maintenance, plus a $150K-$500K internal IT team.

44. KLAS Research. U.S. Hospital EHR Market Share (2024). https://klasresearch.com/

45. Nightingale OS (internal). Market Research & Positioning Report; Vision; Digital Pilot metrics (June 2026).

Prepared by Regulation Loop, Inc. (Nightingale OS). Sections I-IV rest on the public sources above; Section V is interpretive. Forward-looking and pilot figures are modeled projections, not observed outcomes. June 2026.