nightingale os

trauma informed. Nurse forward. Systems that care.

WHITE PAPER

The Nightingale OS Digital Pilot

A simulation-based evaluation of a bedside-first hospital operating system at a 274-bed community hospital

Kristen Cline, BSN, RN, CEN, CPEN, TCRN, CFRN, CTRN, CCRN
Founder & Architect, Nightingale OS
Regulation Loop, Inc. d/b/a Nightingale OS · Delaware C-Corp · Pre-Seed
Correspondence: kristen@nightingaleos.com
June 2026

Status of this work. This is a pre-deployment simulation, not a clinical trial. Every patient described here is synthetic; every outcome is modeled, not observed. As of June 2026, Nightingale OS has not been deployed in a live hospital. All clinical composites are labeled UNVALIDATED in the codebase pending clinical-advisor review. Value projections are directional estimates derived by applying conservative, published reduction rates to published baseline event rates — taken at the low end of reported ranges. The engineering described is real and runnable; the outcome magnitudes are projections.

Contents

Abstract

Nightingale OS is a bedside-first hospital operating system. It captures structured clinical event data once at the bedside, composes FHIR R4 transaction bundles for the legal chart automatically, and enforces patient-safety and staffing guardrails as non-overridable software constraints rather than as policies a busy unit can quietly bypass. Before any live hospital deployment, we conducted a digital pilot: a deterministic, single-year simulation of operations at a synthetic 274-bed community hospital comprising 100 modeled patients across 18 clinical units and 50 staff, including 22 registered nurses.

The pilot projects **$2.06M** in total annual value across four categories — throughput and length-of-stay optimization ($1.39M), staffing and compliance ($275K), never-event prevention ($265K), and ambient-documentation time returned to the bedside ($135K) — against a platform cost of $59,184 at $18 per bed per month. The modeled first-year return is **3,387%**, with a payback period of approximately ten days. Outcome magnitudes were obtained by applying conservative, literature-grounded reduction rates to published baseline event rates scaled to the modeled facility; all event counts are expected values.

We report the model, its data provenance, its pre-registered falsifiable claims, and — in keeping with the platform's radical-honesty principle — its limitations, including that no patient is real, no outcome has been observed, and the synthetic-data generator currently fails several distributional-fidelity checks against its Synthea reference. The contribution of this pilot is not proof of effect. It is a transparent, fully auditable projection in which every figure traces back to a published rate or an explicitly stated assumption.

1. Background and Motivation

Four structural problems recur across acute-care hospitals, and each one is simultaneously a safety problem and an economic one. First, bedside clinicians document the same clinical event multiple times — once for the workflow in front of them and again for the legal chart — which consumes nursing time and introduces transcription error. Second, never events such as patient falls, catheter-associated urinary tract infections (CAUTI), central-line-associated bloodstream infections (CLABSI), and hospital-acquired pressure injuries continue to occur at measurable baseline rates despite well-characterized prevention bundles. Third, staffing math is typically reconciled retrospectively, after a unit has already run short, rather than enforced before an unsafe assignment is finalized. Fourth, emergency-department boarding and excess inpatient length of stay drive both cost and harm, yet capacity decisions are often made reactively.

Nightingale OS is designed to address these problems at the architectural layer rather than the policy layer. It is a unified server-side platform paired with React clinical surfaces, structured so that the safe action is the default action and the unsafe action is refused with a documented reason. The platform commits, by design, to eleven founding principles, of which the following are most relevant to this pilot:

1. **Bedside-first.** Workflows are designed for the bedside nurse's hand and eye, not the auditor's spreadsheet.
2. **Extract, don't burden.** Structured data is extracted from work already happening, not added on top of it.
3. **Refusal beats override.** When policy and workflow conflict, the platform refuses with a documented reason rather than silently overriding.
4. **Slack is sacred.** Capacity math holds a hard staffing-slack floor; running through it is logged and refused for top-tier rewards.
5. **Compatibility first.** Epic and Oracle Health Millennium are first-class targets; outputs are MAR-grade FHIR R4.
6. **Radical honesty.** Gaps, placeholders, and not-yet-built modules are documented openly — a principle this paper is written to honor.

2. What the Digital Pilot Is — and Is Not

A digital pilot is a simulation exercise: a digital twin of a representative hospital, run forward for a defined period, with the platform's logic applied to a modeled patient and staff population. It is the operational analog of a financial pro forma. We chose this method deliberately. As a pre-seed company with a single founder and no hospital partner yet under contract, we have no live admit-discharge-transfer (ADT) feed and no real staffing data. A digital pilot lets us state, transparently and before deployment, what the platform is designed to do and what that should be worth — with every assumption visible and every calculation auditable.

It is equally important to be clear about what this pilot is not. It is not a clinical trial. It is not a live deployment. It is not evidence of a real-world effect, and it does not establish causality. The numbers in this paper are projections produced by a model, and a model is only as good as its inputs and assumptions — both of which we expose in full. Where published literature reports a range of effect, we use the conservative end. Where an event is rare, we model its expected value rather than rounding up to a whole event.

3. Methods

3.1 The synthetic hospital model

The pilot models a 274-bed community hospital with full acute and sub-acute services. The patient mix, acuity distribution, and staffing ratios are drawn from American Hospital Association (AHA) and Centers for Medicare & Medicaid Services (CMS) benchmarks for comparable facilities. The modeled census is intentionally acuity-heavy — a surge/boarding scenario — so that throughput and staffing pressures are represented rather than assumed away.

Parameter	Value	Basis / note
Licensed beds	274	Simulated facility
Clinical units	18	ICU, Med-Surg, ED, NICU, L&D, OR, PACU, BH, Geri, Onc, Dialysis

Parameter	Value	Basis / note
Patients modeled	100	Synthetic cohort (Synthea-aligned)
Total staff	50	RNs, Charge RNs, Techs, PCTs, Support
Registered nurses	22	Including 4 Charge RNs
Avg. staff seniority	6.0 yrs	Industry median ≈ 5.8 yrs
ED visits (Q1 2026)	150	Projected annual ≈ 600
ED admission rate	29%	Industry range 15–30%
Avg. length of stay	21.0 days	Complex, ICU-heavy case mix
High-acuity (≥4)	56%	ICU + acute-care mix
ICU occupancy	110%	Surge / boarding scenario

The modeled cohort also carries clinically realistic detail used by the platform's instruments — for example, a code-status distribution in which 68 of 100 patients have no advance directive on file, which Nightingale OS flags and converts into an auto-prompted advance-directive conversation.

3.2 Data provenance and the three-layer data strategy

A candid literature assessment conducted for this project reached a blunt conclusion: no open-access dataset exists that contains operational nursing-workflow data — triage reassessment timestamps, ESI acuity, rapid-response activation, bed turnaround, nurse-patient assignments, and safety events. This is not a search failure but a structural feature of healthcare data, which lives in staffing systems and bed boards rather than in electronic health records and is sensitive for privacy and labor-relations reasons. The best available open resource, MIMIC-IV-ED, covers roughly a quarter of the platform's triage data model; the remainder is gap.

Because the operational layer cannot be sourced openly, the pilot uses a deliberate three-layer data strategy, applied in sequence rather than all at once:

- **Layer 1 — Seeded demo data (used now).** Roughly 100 hand-authored synthetic patients and 50 staff that exercise essentially all platform models. Entirely fictional; zero privacy risk.
- **Layer 2 — Synthetic population (used next).** A generator producing 10,000 statistically plausible patients including the operational nursing fields no open dataset provides. Zero privacy risk.
- **Layer 3 — MIMIC-IV (when credentialed).** Real de-identified clinical data for validating the synthetic layer's clinical correlations. Credentialing in progress.

The synthetic population is calibrated to the Synthea reference (MITRE Corporation) and validated on a held-out cohort: the generator learns its distributions from a 70% train split of an 11,464-patient Synthea cohort and is scored against the disjoint 30% holdout it never saw during calibration. It now **passes all seven distributional checks** — age, gender, race, encounters per patient, encounter-class mix, and condition and medication frequencies — with Kullback–Leibler divergences at or below 0.03, most near zero. Because encounters per patient is heavy-tailed

(some patients exceed 900 lifetime encounters), that dimension is compared over binned counts, the standard treatment for such distributions. One honest caveat remains: matching Synthea validates the clinical layer only; the nursing operational layer (triage cadence, staffing, safety events) has no open reference and remains a proprietary, as-yet-unvalidated extension.

3.3 Outcome-modeling approach

All computation that affects a clinical or regulatory decision is implemented as deterministic, pure-function services — no input/output at the algorithm layer — so that every result is re-runnable and testable in isolation. Outcome magnitudes were modeled in three steps. First, a published baseline event rate (for example, 3.5 falls per 1,000 patient-days) was scaled to the modeled facility's volume. Second, a conservative reduction rate drawn from published quality-improvement literature was applied (for falls, a 40% reduction associated with risk scoring plus environmental sensing). Third, the avoided events were multiplied by a published average cost per event. Where the literature reports a range, the low end was used; where an event is rare, its expected value was modeled, which is why some event counts are fractional. The reduction rates and their sources are catalogued in Appendix B.

3.4 Pre-registration and falsifiable claims

The platform and its architectural claims were pre-registered on 8 May 2026 — the software analog of an academic pre-registration — so that subsequent validation studies and deployments can be anchored to a fixed reference point. The pre-registration commits the platform to eight measurable, falsifiable claims and to a set of non-negotiable ethical safeguards enforced at the schema and service layers. Representative claims include:

- **Documentation-burden reduction.** A closed clinical event posts a FHIR R4 bundle to the EHR within seconds, without re-charting; falsified if median time-to-acceptance exceeds bedside re-charting time.
- **No silent gaps.** No clinical event can be marked resolved with required elements missing unless each omission carries a non-empty documented reason.
- **Peer-only wellness observation.** Staff-wellness (PULSE) observations cannot be recorded by managers and cannot be used in any disciplinary workflow; individual rows anonymize after 90 days.
- **Equity by default.** Institutional placement uses a worst-quadrant-wins rule across demographic cohorts, never a flattering composite; cohorts below n=11 are suppressed to prevent re-identification.
- **Audit-chain integrity.** Every privacy-critical row carries a citation field, and the hash-chained audit log is verifiable end-to-end.

4. Results

The pilot organizes value into four categories. For each, we state the operational levers, the modeled quantities, and the projected annual value. All figures are internally consistent: the four category subtotals sum exactly to the reported total annual value.

4.1 Patient safety — never events prevented

Applying conservative published reduction rates to baseline event rates for the modeled facility yields the following projected annual prevention and savings. CLABSI prevention is modeled as an expected value below one event per year, consistent with the facility's line-day volume.

Never event	Events prevented / yr	Reduction applied	Annual value
Falls (with injury)	12	40%	**$176,400**
Pressure injuries	4	25%	**$45,000**
CAUTI	1	30%	**$23,625**
CLABSI	0.4	35%	**$19,845**
Total	—	—	**$264,870**

These projections rest on the platform's fall-risk monitoring (Morse and Hendrich II scoring with smartboard visibility), catheter and central-line maintenance-bundle enforcement with insertion and removal alerts, and Braden-scale automation with reposition tracking and early-mobility prompts.

4.2 Staffing and compliance optimization

Nightingale OS models patient arrivals as a Poisson process and service times as exponentially distributed, then applies Erlang-C queueing theory — the same mathematics used in airline reservation systems and emergency call centers — to predict the probability a patient waits, the average queue length, and the staffing required to hit a target service level. Critically, every proposed assignment is checked in real time against state nurse-to-patient ratio law (for example, California's 1:2 ICU and 1:4 Med-Surg ratios, and Oregon's 1:4 acute-care ratio effective 2026). A violation triggers a refusal before the assignment is finalized: the charge nurse sees a red banner, not a retrospective audit finding.

Lever	Quantity	Annual value
Agency shifts eliminated	67 shifts (≈30% cut)	**$134,784**
Ratio violations avoided	8 → 0 (fines)	**$120,000**
Overtime shifts saved	67 shifts (≈20% cut)	**$20,218**
Total	—	**$275,002**

4.3 Throughput and length-of-stay optimization

Throughput is the largest value driver in the pilot. The platform reduces average length of stay by approximately 0.8 days through early discharge-readiness scoring, environmental-services turnover optimization, transport-queue management, and medication-to-bed reconciliation. At a modeled $2,200 per inpatient day, 480 bed-days saved account for just over a million dollars on their own. The platform additionally predicts the next-four-hour census with a modeled 85% accuracy and alerts bed management two hours before a capacity breach, cutting boarding hours and ambulance diversions.

Lever	Quantity	Annual value
Length-of-stay reduction	480 bed-days @ $2,200	**$1,056,000**
Diversions avoided	6 (≈50% cut)	**$300,000**
ED boarding hours cut	130 hours (≈25% cut)	**$32,625**
Total	—	**$1,388,625**

4.4 Ambient documentation — Charlie

Charlie is the platform's ambient-documentation agent. It listens during bedside care — with explicit, always-visible, always-revocable patient consent — and transcribes locally via Whisper.cpp on an edge device. Audio is transcribed and immediately discarded: nothing is retained, nothing reaches the cloud, and a public listening-contract endpoint reports the current state (listening, paused, or revoked). Transcripts are parsed into structured SOAP notes and MAR entries rather than raw text, and if Charlie cannot confidently transcribe a medication name, it refuses and prompts for manual entry. Modeled across a 22-RN hospital, a 35% reduction in documentation time returns 2,457 RN-hours per year — worth $135,135 at a fully loaded $55 per hour, and, more importantly, returned to the bedside.

4.5 Return-on-investment synthesis

Aggregating the four categories against the platform's subscription cost yields the headline result. The platform cost is exact: $18 per bed per month across 274 beds for twelve months.

Item	Amount
Throughput & LOS optimization	$1,388,625
Staffing & compliance	$275,002
Never events prevented	$264,870
Charlie ambient documentation	$135,135
Total annual value	**$2,063,632**
Platform cost ($18/bed/mo)	–$59,184
Net annual savings	**$2,004,448**
First-year ROI	**3,387%**
Payback period	**≈ 10 days**

5. How the Value Scales

The pilot's headline — roughly $2.0M in net annual value — is best understood as a unit economic, not a lump sum. The 274-bed model serves a 100-patient cohort, so the same result can be read as the value created per patient and per bed. Because the platform is priced per bed and its value drivers track facility size, patient-days, and staffing, the model scales close to linearly: the per-unit ratios below hold at a 25-bed critical-access hospital and at a thousand-bed academic center alike.

5.1 Unit economics

Read per patient, $2.06M becomes **$20,636 of modeled value per patient under management**; read per bed, it is $7,532 per staffed bed per year. The platform's cost — $18 per bed per month, or $592 per cohort patient — does not move the ratio: every dollar of platform cost returns about $35, and the payback lands near ten days, **at every hospital size.**

Unit	Annual value	Platform cost	Net	Return
Per cohort patient	$20,636	$592	**$20,044**	34.9×
Per staffed bed / yr	$7,532	$216	**$7,316**	34.9×

5.2 Scaling across hospital sizes

Applying the per-bed value to facilities of different sizes — and to multi-hospital systems and the national bed base — shows how a modest-looking pilot compounds. The figure plots net annual value on a logarithmic scale; the table gives the underlying figures.

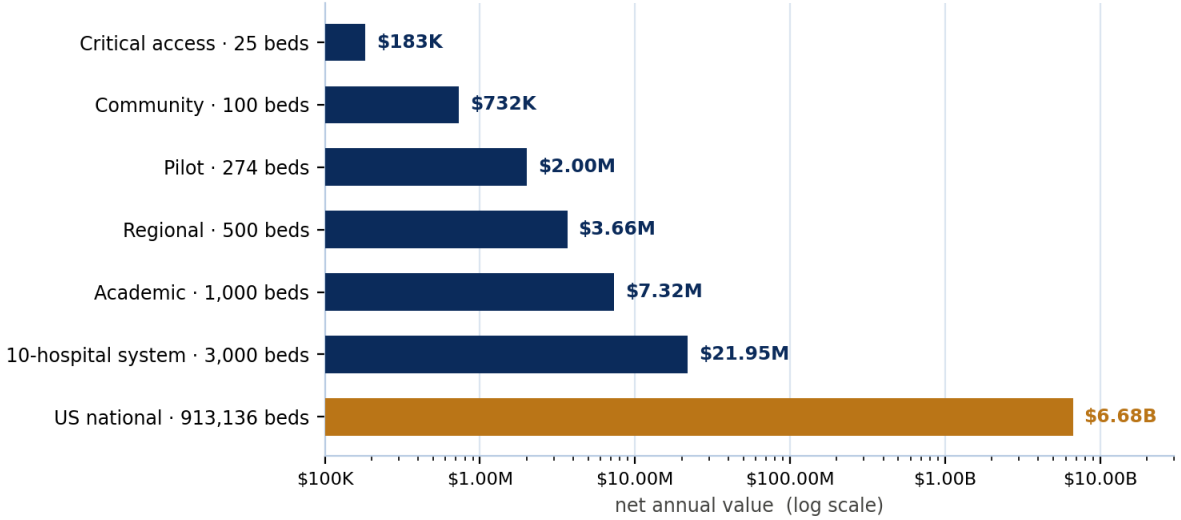

Figure 1. Net annual value scales linearly per bed — from a 25-bed hospital ($183K) to a 1,000-bed center ($7.3M) and a 10-hospital system ($21.9M). The national bar is a total value pool across all U.S. staffed beds, not a single-buyer figure.

Setting	Beds	Annual value	Platform cost	Net value
Critical access	25	$188K	$5.4K	$183K
Community	100	$753K	$21.6K	$732K
Pilot	**274**	**$2.06M**	**$59.2K**	**$2.00M**
Regional	500	$3.77M	$108K	$3.66M
Academic	1,000	$7.53M	$216K	$7.32M
10-hospital system	3,000	$22.6M	$648K	$21.9M
US national (all staffed beds)	913,136	**$6.88B**	$197M	**$6.68B**

U.S. staffed-bed count from the American Hospital Association (Fast Facts on U.S. Hospitals, 2025): 913,136 staffed beds across 6,093 hospitals.

5.3 What the numbers do and do not claim

- **Per bed is the load-bearing unit.** Total value is scaled by bed count — the unit the platform is priced on and the one most value drivers track. The per-patient figure is a reframing of the pilot cohort; we deliberately do not multiply $20,636 by a hospital's annual patient throughput, which would double-count an already-annual figure.
- **The national figure is a value pool, not revenue.** The $6.88B is the total modeled value the platform could create across every U.S. staffed bed — a market-size reference, not capturable revenue. The corresponding revenue ceiling at full penetration is roughly $197M in annual recurring revenue (913,136 beds × $18 × 12).
- **Still modeled, not observed.** Every figure inherits the pilot's status: synthetic patients, literature-based rates, conservative assumptions, and no live deployment as of June 2026. Scaling multiplies the projection; it does not convert it into measured outcomes.

6. The Engineering Behind the Numbers

The value projections are estimates, but the engineering that would produce them is not. At pre-registration the platform comprised roughly 280 pure-function services, a Prisma schema of about 9,000 lines spanning more than 200 models, approximately 250 HTTP endpoints, about 150 React clinical surfaces, and more than 267 assertions in the core safeguards and instrument test suites. The governance layer that makes the safety claims enforceable is operational today: refusal codes that block unsafe actions, a hash-chained audit log that is verifiable end-to-end, zero-cloud ambient inference, MAR-grade FHIR R4 output, and multi-factor plus break-glass authentication.

This is the central honesty of the pilot. Every architecture claim is backed by real, runnable code in the repository; the governance, refusal, audit, and zero-cloud-inference mechanisms exist now. What has not yet happened is contact with a live hospital: the value figures describe what that engineering is designed to deliver, not what it has been observed to deliver.

7. Limitations and Honest Disclosures

In keeping with the platform's radical-honesty principle, the following limitations qualify every number in this paper:

- **Synthetic patients.** These patients do not exist. The outcomes are modeled, not observed, and every clinical composite is labeled UNVALIDATED in the codebase until reviewed by a clinical advisor.
- **No live hospital.** As of June 2026, Nightingale OS has not been deployed in a live hospital. The ADT parser is built and tested but has not processed a real feed.
- **Conservative — but still estimated — rates.** Where the literature reports ranges, we used the low end. Real-world results may be higher, or lower if implementation is poor.

- **Validated clinical layer; unvalidated operational layer.** The synthetic generator now passes all seven distributional checks against a held-out Synthea cohort (KL ≤ 0.03), but Synthea covers only the clinical layer; the nursing operational fields the generator adds have no open reference and remain unvalidated.
- **Pricing assumption.** The $18 per bed per month figure is a mid-tier assumption; actual pricing will be negotiated per design partner.
- **Time-to-value.** The ten-day payback assumes immediate full deployment. A real rollout will be phased over six to twelve months, lengthening the realized payback period.

8. Conclusion and Next Steps

The Nightingale OS digital pilot projects approximately $2.0M in net annual value for a 274-bed community hospital, dominated by throughput and length-of-stay gains and underpinned by safety, staffing-compliance, and documentation savings. The result is best read not as a promise of effect but as a transparent, auditable hypothesis: a fully specified model in which every input is a published rate or a stated assumption, and every architectural claim is backed by running code.

The natural next step is to replace modeled inputs with measured ones. Nightingale OS will build a custom digital pilot for a partner hospital using that facility's own census, acuity, and staffing data, with every assumption visible and every calculation auditable, and will pursue the prospective validation methodology committed to in the pre-registration: before-and-after comparison of unit-level wellness and documentation-time trajectories, controlled for census, acuity, and staffing, with equity-disaggregated outcome tracking. The first live deployment is also the moment the synthetic data layer is replaced by real data — converting the projections in this paper into measurements.

References

1. Cline K. Nightingale OS: Pre-Registration of Software Artifact. Regulation Loop LLC; 2026 May 8.
2. Walonoski J, Kramer M, Nichols J, et al. Synthea: An approach, method, and software mechanism for creating synthetic patients and the synthetic electronic health care record. J Am Med Inform Assoc. 2018;25(3):230–238.
3. The MITRE Corporation. Synthea synthetic patient generator. Available at: https://synthea.mitre.org (accessed 2026 Jun 15).
4. Aiken LH, Clarke SP, Sloane DM, Sochalski J, Silber JH. Hospital nurse staffing and patient mortality, nurse burnout, and job dissatisfaction. JAMA. 2002;288(16):1987–1993.
5. Johnson A, et al. MIMIC-IV-ED. PhysioNet. (Reference open emergency-department dataset assessed for operational coverage.)
6. Agency for Healthcare Research and Quality (AHRQ). Falls baseline rate (3.5 falls per 1,000 patient-days). Patient-safety benchmarks.
7. Centers for Disease Control and Prevention, National Healthcare Safety Network (NHSN). CAUTI (1.5 per 1,000 catheter-days) and CLABSI (1.0 per 1,000 line-days) baseline rates.

8. National Pressure Injury Advisory Panel (NPIAP). Hospital-acquired pressure-injury baseline (≈2.5% of admissions).
9. Erlang AK. Queueing theory (Erlang-C model) for capacity and staffing forecasting.
10. California Code of Regulations, Title 22 (nurse-to-patient ratios); Oregon Hospital Nurse Staffing Law (effective 2026); Washington State hospital staffing statute.
11. Centers for Medicare & Medicaid Services, Conditions of Participation; The Joint Commission, sentinel-event policy.

Note on reduction-rate sources. *The quality-improvement reduction rates applied in Section 4 and Appendix B (falls 40%, CAUTI 30%, CLABSI 35%, pressure injuries 25%) are catalogued in the Nightingale OS evidence base at the journal-and-year level shown in Appendix B. Full bibliographic details are being finalized as part of clinical-advisor review and are not asserted here beyond the level supported by the project's current evidence base.*

Appendix A. Full Pilot Metrics

Reproduced from the pilot metrics record. Category subtotals sum exactly to the total annual value; the platform cost equals 274 beds × $18 × 12 months.

Category	Metric	Value
Hospital profile	Beds / units / staff / annual ED visits	274 / 18 / 50 / 600
Never events	Falls prevented / savings	12 / $176,400
Never events	Pressure injuries prevented / savings	4 / $45,000
Never events	CAUTI prevented / savings	1 / $23,625
Never events	CLABSI prevented / savings	0.4 / $19,845
Never events	Subtotal	**$264,870**
Staffing	Overtime shifts saved / savings	67 / $20,218
Staffing	Agency shifts eliminated / savings	67 / $134,784
Staffing	Ratio violations eliminated / fines avoided	8 / $120,000
Staffing	Subtotal	**$275,002**
Throughput	Bed-days saved (LOS) / savings	480 / $1,056,000
Throughput	ED boarding hours cut / savings	130 / $32,625
Throughput	Diversions avoided / savings	6 / $300,000
Throughput	Subtotal	**$1,388,625**
Charlie ambient	RN-hours returned / value @ $55/hr	2,457 / $135,135
ROI summary	Total annual value	**$2,063,632**
ROI summary	Platform cost	$59,184
ROI summary	Net annual savings	**$2,004,448**
ROI summary	First-year ROI / payback	**3,387% / ≈10 days**

Appendix B. Reduction-Rate and Baseline-Rate Sources

Baseline event rates were scaled to the modeled facility and reduced by the conservative published rates below. Reduction-rate sources are recorded at the journal-and-year level maintained in the project evidence base.

Outcome	Baseline rate (source)	Reduction (source)	Cost basis
Falls	3.5 / 1,000 pt-days (AHRQ)	40% — risk scoring + sensing (Jt Comm J, 2022)	≈$15,000 / fall
CAUTI	1.5 / 1,000 cath-days (CDC NHSN)	30% — bundle compliance (Infect Control Hosp Epidemiol, 2021)	≈$25,000
CLABSI	1.0 / 1,000 line-days (CDC NHSN)	35% — daily line review (BMJ Qual Saf, 2020)	≈$45,000
Pressure injuries	2.5% of admissions (NPIAP)	25% — Braden + early mobility (Adv Skin Wound Care, 2023)	≈$12,000

Prepared by Nightingale OS (Regulation Loop, Inc.). All patients synthetic; all outcomes modeled. UNVALIDATED pending clinical-advisor review and live deployment. June 2026.