WHITE PAPER · SEPTEMBER 2026

One patient record across every line your practice runs.

MedOp is the AI-native system of record for behavioral health, rehab therapy (PT/OT/SLP) and medical aesthetics — built vertical-deep, priced per provider with everything included, and run by a governed workforce of 27 AI agents.

MedOp · medop.io

Platform build v7.63.0 · Document version 1.1 · 2 September 2026

Rachael Pereira, Chief Executive Officer · Paul Pereira, Lead Development Engineer

Contents

1. Executive summary

Independent behavioral health, rehab therapy and aesthetics practices run on the same thing: people doing administrative work that the software should have done. Notes written after hours, prior authorizations chased by phone, claims denied for reasons that were knowable before the visit, patients who never came back because nobody called. The vendors serving these practices sell generic charting with the specialty's name on the template, meter every add-on, mark up card payments, and charge to export the practice's own data.

MedOp replaces that stack. It is a single system of record — scheduling, intake, clinical documentation, coding, claims and revenue cycle, e-prescribing, a patient portal with secure messaging, payments and reporting — with a workforce of 27 specialised AI agents underneath it that draft, triage and propose while a human approves. It is built deep for three verticals rather than wide for thirty, and a practice can enable one line or all three on **one patient record** at the same price. No vendor in a 21-company competitive scan can do that.

27	3	$0	3,700+
AI agents in four pods, on a 37-type governed registry	verticals, composable on one patient record	metered fees, payment markup, onboarding, migration or export	automated tests behind every release (v7.63.0)

The paper makes five claims and supports each with evidence:

1. **The problem is labour, not software features.** Physicians and their staff spend 13 hours per physician per week on prior authorization alone (AMA, 2026); 11.8% of claims are denied on first submission (Kodiak, 2025); 86% of rehab therapists report burnout and 59% name documentation as the primary driver (Ensora, 2026).
2. **The three verticals are large, growing and structurally under-served.** 61.5 million US adults had a mental illness in 2024 and only half were treated (SAMHSA); offices of PT/OT/SLP and audiologists employ 457,000 people (BLS); 10,488 medical spas averaged $1.40 million in revenue in 2023 (AmSpa). None of the three is gated by federal EHR certification, and integrated practices that run more than one line have no vendor at all.
3. **MedOp is a system of record, not an AI overlay.** The overlay category is now a commodity — the incumbents bundle a scribe for free — and is contractually terminable by the EHR it sits on. MedOp migrates the practice in and owns the record. Its first production migration moved 2,868 patients, 18,037 appointments and 11,468 active medications.
4. **The differentiator is the machinery, not the agent count.** Psychotherapy CPT minute bands enforced from the session clock, the Medicare 8-minute rule computed from treatment minutes, a Good Faith Exam gate that refuses to chart an injectable without one, 42 CFR Part 2 disclosure accounting, and a Collaborative Care registry — every one of them undocumented by the incumbents — on an agent substrate with calibrated confidence, approval gates, per-agent kill switches and a tamper-evident PHI audit chain.
5. **The commercial model restores expansion revenue and removes the reasons practices distrust their vendor.** Published per-provider pricing ($129–$229), an AI Workforce meter billed per completed unit of work with a hard cap the practice sets, a 0% platform fee on payments (MedOp takes no percentage — the practice keeps its own merchant account and rate), and a seven-clause Fair Terms Charter written into the subscription agreement.

> **Why now**
>
> Ambient documentation became a $600 million category in 2025 (Menlo Ventures) and is now given away inside the incumbents' EHRs. The value has moved from "an AI that writes the note" to "a system that runs the practice and is accountable for the outcome." The companies capturing that value — Candid Health at 180% net dollar retention, Adonis above 130% — are priced on expansion. MedOp is built and priced the same way, for three verticals the

> horizontal players serve badly.

2. The problem: practices run on labour and leak revenue

A small practice is an administrative machine with a clinic attached. The machine is staffed by people whose day is consumed by documentation, phone calls, payer portals and rework, and the cost of that labour is rising faster than reimbursement. The evidence is consistent across specialties.

2.1 Documentation and burnout

Finding	Figure	Source
Physicians reporting at least one burnout symptom (2025)	41.9%	AMA Organizational Biopsy, Apr 2026
Physicians spending more than 8 hours a week on the EHR outside clinic hours	20.9%	AMA, Aug 2024
EHR time per 30-minute primary-care visit, including after-hours work	36.2 min	JAMA Network Open via AMA, Jan 2024
Rehab therapists (PT/OT/SLP) reporting burnout; naming documentation as the primary driver	86% / 59%	Ensora Health, Apr 2026 (n>500)
Psychologists citing administrative burden (pre-authorization, audits) as a barrier to care	62%	APA Practitioner Pulse Survey, 2024
Physicians who now use AI professionally; using it for documentation or coding	81% / 28%	AMA AI Survey, Mar 2026

2.2 Revenue leakage

Finding	Figure	Source
Initial claim denial rate, 2024 (up 2.4% year on year)	11.81%	Kodiak Solutions, May 2025 (2,100+ hospitals, 300k physicians)
Cost to rework one denied claim, 2023 (from $43.84 in 2022)	$57.23	Premier Inc., Feb 2025
Share of adjudication spend that is “potentially unnecessary”	$18B of $25.7B	Premier Inc., Feb 2025
Providers with denial rates of 10% or more; saying clean claims are harder than a year ago	41% / 68%	Experian Health State of Claims, Sep 2025
Denied in-network marketplace claims that were ever appealed	<1%	KFF, Mar 2026
Prior authorizations per physician per week; hours per week spent on them	39 / 13 hrs	AMA Prior Authorization Survey, May 2026 (n=1,000)
Physicians with staff who work exclusively on prior authorization	40%	AMA, May 2026

These are not hospital problems that skip small practices. They fall hardest on practices without a revenue-cycle department, where the person chasing the payer is also the person answering the phone. The Experian survey found 43% of provider organisations understaffed in the revenue cycle and only 14% using AI for claims — of whom 69% saw fewer denials.

2.3 The software the market sells them

The systems these practices buy were designed as digital filing cabinets with billing attached, and their commercial models have hardened around the practice's inability to leave. In a review of public terms and customer reviews conducted for MedOp's positioning work (September 2026), the pattern was consistent: subscription prices that are only 7–25% of what a practice actually pays once claims, statements, texts, e-prescribing, telehealth and payment processing are metered on top; card-processing rates of roughly 3.1% plus $0.30 against an interchange-plus cost near 1.8–2.3%; no obligation to return the practice's data on termination; price increases at the vendor's election with no cap; and, in one aesthetics platform customer's public review, a fee of over $1,000 to export the practice's own patient records.

The result is a market where clinicians rate their tools poorly — only 18% of physicians report a "strong" or "elite" EHR experience (KLAS Arch Collaborative, 2025) — but rarely switch, because switching is engineered to be painful.

3. The market: three verticals, one under-served segment

MedOp is deliberately vertical. The three specialties it serves were chosen because they are large, because their needs are specific enough that general-purpose EHRs serve them badly, and because they sit outside the federal certification and MIPS machinery that shapes the rest of the ambulatory market.

3.1 Behavioral health

Demand exceeds supply by a wide margin and the gap is widening. 61.5 million US adults had any mental illness in 2024 and only 52.1% received treatment (SAMHSA NSDUH, 2025). 137 million Americans — 40% of the population — live in a Mental Health Professional Shortage Area, and HRSA projects 2038 shortfalls of roughly 100,000 psychologists and 100,000 mental health counselors (HRSA, Dec 2025). Offices of mental-health practitioners generated $35.7 billion in 2025 (IBISWorld). The specialty's own quality movement is under-tooled: fewer than 20% of practitioners use measurement-based care (Lewis et al., JAMA Psychiatry), and only 21% of therapists surveyed have a reliable way to monitor client outcomes (Ensora, Aug 2026, n=1,124). The number of providers billing Collaborative Care in Original Medicare grew 5.8× between 2018 and 2022, yet the programme reaches only 32–65 patients per 100,000 (Milliman/MMHPI, 2025) — a fast-growing programme that no mainstream behavioral-health EHR documents support for.

3.2 Rehab therapy — PT, OT and SLP

Offices of physical, occupational and speech therapists and audiologists employed 457,060 people as of May 2023, including 87,190 physical therapists, 41,050 occupational therapists and 41,990 speech-language pathologists (BLS OEWS), and generated $53.2 billion in 2025 (IBISWorld). Outpatient PT runs a 9.5% vacancy rate, twice the national average (APTA, 2024). Therapy billing is governed by mechanics that a general EHR does not compute — the Medicare 8-minute rule, the annual KX-modifier threshold ($2,480 per discipline bucket in 2026), CQ/CO assistant modifiers, plan-of-care certification windows and the tenth-visit progress-report requirement. Today those are applied by billers, by hand, from memory.

3.3 Medical spas and aesthetics

A fast-growing, cash-pay category. The United States had 10,488 medical spas in 2023, up from 8,899 the year before, with average revenue of $1.40 million per location, 245 patient visits a month, $527 per visit and 81% single-location ownership (AmSpa State of the Industry, 2024). Baird sizes the category at roughly $15 billion growing about 15% a year. It is cash-pay, so the economic lever is payments rather than claims — and it is regulated at the state level through Good Faith Exams, supervision rules and consent, which the leading aesthetics platforms outsource to point solutions at $27.99 per exam.

3.4 The segment nobody serves: integrated practices

> **The composition problem**
>
> A behavioral-health practice that also runs a medical spa — MedOp's own pilot customer is one — operates two systems today with two records for the same person, because every vendor is built for one line of business. The same is true of a PT clinic with an aesthetics room, or an integrative practice offering therapy and wellness services. In MedOp a practice ticks one line or all three; the encounter decides which coding rules, code sets, compliance gates, templates and agent behaviours apply, and the patient record stays one record. Of the 21 vendors in MedOp's September 2026 competitive scan, none can do this at any price.

3.5 Why these verticals avoid the regulatory drag

ONC health-IT certification is voluntary. It matters where it gates MIPS Promoting Interoperability, and MIPS eligibility requires a clinician to exceed all three of $90,000 in Part B allowed charges, 200 beneficiaries and 200 covered services. Cash-pay aesthetics, cash-pay behavioral health and most PT/OT/SLP practices sit outside those thresholds. Building for them means MedOp can invest in specialty depth rather than in a certification programme that the HTI-5 proposed rule (90 FR 60970) would shrink by more than half — and it keeps MedOp outside the information-blocking "actor" category that attaches only to developers of certified modules.

3.6 Software spend

The US practice-management software market was $5.89 billion in 2024 and is projected to reach $13.70 billion by 2033, a 10.0% compound annual growth rate, with integrated products holding 75% share (Grand View Research, 2025). The US behavioral-health software and services market alone is projected to double from $1.49 billion in 2024 to $2.99 billion by 2030 (Grand View Research). Healthcare AI spending tripled to $1.4 billion in 2025, and coding and billing automation was already a $450 million line (Menlo Ventures, Oct 2025). Digital-health venture funding reached $14.2 billion in 2025, with AI-enabled companies taking 54% of the dollars (Rock Health, Jan 2026).

4. The platform: scope of the system of record

MedOp is one application, one database and one patient record. Everything a practice needs to see a patient, document the visit, get paid and keep the patient coming back runs inside it. Nothing below is an add-on with its own price.

Scheduling and front office

Multi-provider calendar with real availability computation, no-show risk, waitlist and reactivation slot ranking, visit reasons tied to vertical rules, and an AI receptionist that answers the phone and books into the same calendar.

Intake and onboarding

Digital intake forms, insurance capture and pre-visit eligibility integrity checks, self-serve CSV migration for patients, appointments and coverage, and a Spanish-language patient journey built on validated translations.

Clinical documentation

Ambient recording with asynchronous record-to-chart, specialty templates, structured outcome measures (PHQ-9, GAD-7, PHQ-2, C-SSRS crisis alerting), therapy worksheets, and note intelligence that summarises, extracts actions and drafts patient instructions.

Coding and compliance

Grounded ICD-10 and CPT suggestion against a 98,000-code catalog, documentation-adequacy review against billed codes, NCCI PTP and MUE edit tables, time-based psychotherapy coding, and a real-time HIPAA anomaly monitor on PHI access.

Claims and revenue cycle

Charge capture, denial-risk scoring before submission, prior-authorization determination and tracking, denial appeal and underpayment-dispute drafting for human signature, an RCM command centre, payer contracts and rates, and 14- and 90-day cash-flow forecasting.

e-Prescribing and safety

Electronic prescribing with EPCS, TOTP-gated controlled-substance workflows, drug–drug interaction and allergy checks, and a refill queue.

Patient portal, messaging and payments

Passwordless patient portal with two-way secure messaging, results release policy, proxy access, and patient payments on Stripe Connect Standard so the practice owns its merchant account and rate.

Engagement and reputation

TCPA-compliant reminders and recall sequences with quiet-hours enforcement, campaign intelligence with A/B arms, patient lifetime-value and churn scoring, and HIPAA-safe replies to public reviews.

Telehealth

Driver-based telehealth with a BAA-covered vendor path and magic-link patient joining.

Reporting and the programmable surface

Curated, audited reports with CSV export; a public REST API with OpenAPI, an MCP server for agent-to-agent integration, signed webhooks and a developer sandbox; FHIR R4 read import.

4.1 Composable verticals

Each vertical contributes its own code sets, rule packs, templates and agent behaviours. A practice enables the lines it runs; the price does not change. The table shows the signature capability in each line — the one that is visible in a sixty-second demonstration and that MedOp's competitive scan found undocumented by every incumbent in that vertical.

Vertical	Signature mechanics	Why it matters	Status
Behavioral health	Encounter timer driving CPT minute-band enforcement (90832 / 90834 / 90837) with an audit-defence path for 90837; 42 CFR Part 2 consent-scoped disclosure and disclosure accounting aligned to the 2024 Final Rule; Collaborative Care registry and care-manager time log (99492/99493/99494); measurement-based care with value-based coverage reporting	90837 is the highest-audit-risk code in the specialty; CoCM and Part 2 are undocumented by all six BH incumbents scanned	Live pilot; minute bands and CoCM registry in build
PT / OT / SLP	The 8-minute rule as a pure function, property-tested over every minute from 0 to 180; KX threshold ledger with auto-attach at $2,480 and a hard alert at $3,000; CQ/CO derived from who performed the minutes; plan-of-care certification windows and tenth-visit progress-report gating; MPPR in payment forecasting	Removes the biller from the loop on every timed code; a 45-minute session charts in under a minute and bills correctly	Specified; scheduled after behavioral health and aesthetics
MedSpa & aesthetics	Good Faith Exam and supervision attestation that hard-block charting of an injectable when unsatisfied; versioned consent pinned to the encounter with a separately revocable photo release; lot, expiry and witnessed wastage on every administration; memberships, packages, gift cards and injector commission; adverse-event export	No aesthetics platform ships an enforced native GFE — a $27.99-per-exam point-solution category exists to fill the gap	GFE gate in the Proof Scope; inventory and commerce to follow

4.2 Cross-line containment

Composition is only safe to sell if the lines are sealed where the law requires it. 42 CFR Part 2 data created in the behavioral line is kept out of an aesthetics encounter's agent context by a retrieval-scope rule enforced in the agent runner — never by a prompt — with line-scoped consent and revocation. The containment is proven by test, and those tests are part of the release gate.

4.3 Migration in, export out

MedOp does not sync with the incumbent forever; it migrates the practice in and becomes the record. The migration adapter is one-way, run-once and resumable with a written manifest; the first production migration moved 2,868 patients, 18,037 appointments and 11,468 active medications from a legacy PM/EHR. Self-serve CSV import now covers patients, appointments and insurance. The reverse direction is a contractual promise: a one-click full export — clinical, financial, documents and images in CCDA, JSON, CSV and original binaries — warranted to complete within one business day, with a service credit if it does not. It lives in the practice's main settings, not behind a support ticket.

5. The AI workforce: 27 agents on a governed substrate

Every agent in MedOp owns one job, reports its confidence on every action, and commits nothing to the record without either a calibrated confidence above the practice's threshold or a human approval. The agents are operator-facing by design: no agent is exposed directly to patients, which bounds both liability and cost.

5.1 The four pods

Pod	Agents	What they do
Clinical (8)	Ambient Scribe · Note Intelligence · Care Gap Hunter · Risk Scoring · eRx Safety · Compliance Monitor · Compliance Audit · Quality / MIPS	Turn the encounter into a defensible note and code; close quality gaps; flag interactions and allergies; watch PHI access in real time; check documentation against billed codes
Revenue (8)	Charge Capture · Grounded Coding · Denial Prevention · Prior Auth · Revenue Cycle · Denial Appeal · Underpayment Recovery · Financial Forecasting	Propose the charges, ground the codes, score denial risk before submission, determine and draft prior authorizations, triage claims through their lifecycle, draft appeals and underpayment disputes for signature, and forecast cash
Front Office (4)	Smart Scheduling · Intake Intelligence · Staff Workflow · Referral Manager	Predict no-shows and rank reactivation slots, verify demographics and coverage before the visit, watch note-completion and burnout signals, and assemble referral packages
Engagement (6)	Patient Outreach · Campaign Intelligence · Review Replies · Review Insights · Patient LTV · Recall Engine	Build the daily recall worklist, run TCPA-compliant reminder and reactivation sequences, draft HIPAA-safe review replies, cluster sentiment, and score lifetime value and churn
Orchestration (1)	Orchestrator	Plans multi-agent workflows as directed graphs and hands each step to the right specialist

Beneath the 27 named agents the registry holds 37 agent types, including a front-line dispatcher, internal document-triage and refill-queue automations, and two meta-agents — a reflection critic that self-reviews every structured output against schema, safety and grounding, and a feedback critic that diagnoses why a rejected output failed. The registry is enforced by a reflection test in CI: the catalog the user sees and the code that runs cannot drift apart.

5.2 The substrate — what makes an agent trustworthy

Prompts are commodity; any competitor can write them in a week. What is hard to replicate is the machinery every agent runs on. MedOp's substrate took eight foundation sprints to build and, in the company's estimate, three to six months for a competent team to reproduce.

Memory and retrieval

Four-tier agent memory with outcome-closing decision logs, and practice-isolated retrieval over pgvector — payer contracts, standard operating procedures and the practice's own won appeals — scoped by practice on every query path.

Reflection and calibration

Every structured output is self-critiqued before it is shown; confidence is rescaled against the practice's own accepted and rejected history so that "0.92" means the same thing next month as it did last month.

Confidence-band escalation

Actions route by calibrated confidence: draft-only by default, human approval in the middle band, and auto-routing only where the practice has opted in above 0.95 and the clinical-safety override is satisfied.

One approval inbox

Every proposal that needs a human lands in one inbox, stamped at creation with priority, SLA and target role, with the supporting context alongside. Approvals can be inherited by later steps of the same workflow inside a five-minute window, and every inheritance is a new audit row.

Kill switches and budgets

Per-practice and global toggles for every agent; a global draft-only switch; a daily autonomy budget that degrades to a cheaper model at 80% and stops at 100%.

Cost attribution in CI

Every model call passes through one cost-attributed chokepoint; a CI gate fails the build on any unattributed call. Per-practice cost of goods is a first-class metric, not an estimate.

Closed-loop improvement

Rejected outputs feed a critique agent; prompt candidates are shadow-executed on a 5% production sample and promoted only when they clear the eval gate. Pre-merge evals and a nightly drift check are part of the release process.

Cost–quality cascade

Non-safety agents run on the fastest model tier first and escalate to the standard or frontier tier on calibrated confidence, not raw confidence — a 25–50% reduction in cascade-eligible call cost at constant acceptance rate.

5.3 How a workflow runs: a denied claim

1. A claim flips to denied. The event lands in MedOp's outbox and triggers the denial-appeal signature workflow.
2. The Orchestrator plans a graph — load the claim, then analyse the denial and validate compliance in parallel, then draft the appeal.
3. The drafting step recalls prior outcomes on this payer and code, retrieves from the won-appeal and payer-contract corpora, pulls in negative examples the practice previously rejected, drafts with the appeal agent, and self-critiques.
4. Calibrated confidence decides the route. Below threshold, the letter queues in the approval inbox with its evidence; a reviewer signs in one click.
5. On approval a deterministic write sends the letter and updates the claim. Telemetry captures tokens, cost, retries, the calibration delta, the gate action and every memory and chunk the agent used.

The same pattern — event, plan, retrieve, draft, reflect, gate, commit, audit — runs prior authorizations, underpayment disputes, charge proposals, recall campaigns and review replies.

5.4 What this is not

It is not a chat window over a general model. A general model does not know this practice's payer contracts, its past won appeals, or which of its drafts staff have rejected; it does not learn, self-critique, route to a human at low confidence, or write an audit row. Every agent in MedOp does all of those, inside a tenant boundary the database itself enforces.

6. What sets MedOp apart

Every incumbent now ships a scribe, several ship a receptionist, and one aesthetics platform markets eleven agents. MedOp does not compete on agent count. It competes on six things that are checkable.

#	Differentiator	What a buyer can verify
1	One patient record across every line the practice runs	Tick a second vertical in settings; the product reconfigures live with the same patients intact. A therapy visit and a neurotoxin visit on the same day follow different rules on one chart.

#	Differentiator	What a buyer can verify
2	System of record, not an overlay	Charting, scheduling, billing, e-prescribing and the portal all run in MedOp; the practice stops paying for a separate EHR. No dependency on an incumbent's API terms.
3	Mechanics the incumbents do not document	Minute-band coding from the session clock; the 8-minute rule as a tested function; an enforced Good Faith Exam gate; Part 2 disclosure accounting; a CoCM registry.
4	A governed agent substrate	Calibrated confidence, one approval inbox, per-agent kill switches, cost attribution enforced in CI, closed-loop improvement — inspectable in the admin console.
5	Tenancy and audit enforced by the database	Postgres Row-Level Security forced on 71 tables under a non-bypass role; a hash-chained, append-only PHI audit log with daily seals; three principal classes with disjoint signing secrets.
6	Fair terms as contract language	Month-to-month always; a warranted one-click export with an SLA and a credit; zero price escalator with published price history; everything in the number; payments at cost; no gag clause; no licence over practice or patient data.

6.1 The competitive landscape in one paragraph

In behavioral health, the incumbents (SimplePractice, TherapyNotes, Valant and others) offer month-to-month terms but generic charting, metered add-ons and no documented support for Collaborative Care or Part 2 accounting. In therapy, WebPT and Jane are strong on scheduling and documentation but leave the 8-minute rule and KX tracking to the biller, and WebPT's agreement permits uncapped price increases. In aesthetics, Zenoti launched an eleven-agent workforce in April 2026 and Boulevard added e-prescribing, but neither ships an enforced Good Faith Exam or witnessed wastage, and the category's transparent software pricing tops out near $328 a month while payments are resold at a markup. Horizontal EHRs — athenahealth, Elation, legacy PM/EHR systems — now bundle ambient documentation at no charge, which is exactly why an AI overlay is no longer a business. None of twenty-one vendors reviewed publishes AI accuracy metrics or an evaluation methodology; MedOp will.

7. Benefits: time, revenue and risk

MedOp is sold as replacing software and labour, and the benefits divide the same way: hours the practice no longer pays for, dollars it no longer loses, and exposure it no longer carries.

7.1 For clinicians

- **The note writes itself, and the code grounds itself.** Ambient capture and asynchronous record-to-chart remove documentation from the evening. athenahealth, the largest deployment to publish a figure, reports roughly six minutes saved per encounter — more than two working days a month per clinician — from native ambient documentation (Aug 2026).
- **Every clinical agent is unlimited and never metered.** A per-note price makes clinicians ration the scribe; MedOp deliberately does not charge one.
- **Specialty-correct by construction.** A 90837 that survives audit on the note alone; a timed therapy session that bills correctly without a biller's touch; an injectable that cannot be charted without its Good Faith Exam.

7.2 For the front office and billing

- **Work arrives done, not assigned.** Prior authorizations determined and drafted, denials scored before submission, appeals researched and packaged, recall lists built with talking points, inbound calls answered and booked.
- **One inbox, one decision.** Staff approve rather than produce. Priority, SLA and evidence are attached to every item.
- **A meter each provider can read.** Units completed, amount billed, and beside it the staff hours and cost avoided — "$126 instead of $840" — so the arithmetic is the practice's, not the vendor's.

7.3 For the owner

- **One bill, one number.** Claims, eligibility, statements, reminders, telehealth, e-prescribing, EPCS, forms, portal, fax, additional locations, onboarding, migration and export are $0, permanently.
- **Payments at cost.** The practice owns its merchant account and rate; MedOp takes 0%. For a medical spa, the payments markup is often the largest line on the software bill.
- **Revenue that stops leaking.** Pre-visit coverage checks, denial prevention, prior-auth tracking and appeal drafting attack the 11.8% first-pass denial rate and the $57 rework cost at the points where they originate.
- **No lock-in.** The data leaves in one click, the price never rises for the life of the subscription, and the contract is month-to-month.

7.4 An illustrative value model

The figures below are an illustration for a three-provider behavioral-health practice, built from published labour and industry benchmarks rather than from MedOp customer data; MedOp will publish measured per-agent scorecards as its founding cohort matures. Assumptions are shown so the reader can substitute their own.

Workload (per week)	Assumption	Hours reclaimed
Documentation on ~90 visits	6 minutes per visit removed from after-hours charting (athenahealth reports ~6 min/encounter)	9.0
Inbound calls, ~150	6 minutes each, answered and booked by the AI receptionist	15.0
Prior authorizations, ~10	25 minutes each drafted and tracked (AMA: 13 hrs/week across physician and staff)	4.2
Denial rework, ~8 claims	35 minutes each researched and packaged (Premier: $57.23 per rework)	4.7
Recall and reactivation outreach	Daily worklist and sequences generated	3.0
Total	Roughly nine-tenths of a full-time front-office equivalent	≈ 36 hrs

At the BLS May 2025 median wage for medical secretaries ($38,290, or $18.41 an hour — about $22 an hour with a 20% load for taxes and benefits), 36 hours a week is roughly $3,400 a month of labour. The same practice's MedOp bill on the metered model is about $1,000–1,200 a month all in, depending on volume — against roughly $1,574 a month it pays today across a behavioral-health EHR subscription, metered add-ons and payment processing. The practice replaces software and most of a salary with a single line item that is smaller than the software alone was.

8. Revenue model and efficiency economics

MedOp's pricing is published, per provider, all-inclusive, and split where the value divides: a platform fee for the system of record, and a variable line — chosen by the customer — for labour the agents perform. Infrastructure is never metered; labour is.

8.1 The platform fee

Providers	Per provider per month	Included
1	$229	Everything: all 27 agents (clinical agents unlimited), claims, eligibility, statements, reminders, telehealth, forms, portal, e-prescribing and EPCS, fax, secure messaging, additional locations, onboarding, migration and export, signed BAA, audit trail, encrypted backups, per-agent kill switches
2 – 9	$189	Same
10 – 24	$159	Same
25+	$129	Same

8.2 The variable line: AI Workforce metering (default)

Class 1 agents — everything a clinician touches inside the visit — are unlimited. Class 2 agents do work a practice would otherwise pay a person to do, and bill per completed unit of work, in dollars, never per token. $75 per provider per month is included; volume bands apply above it; and the practice sets a hard cap at which the agent pauses rather than billing on.

Unit of work	Price	Human equivalent (front-office time at BLS median wage)
Inbound call answered and handled	$0.95	5–8 minutes ≈ $1.80–2.90 at $22/hr
Outbound call — recall, no-show recovery, results	$0.75	Same basis
Payer call — benefits, claim status, escalation	$2.00	12–20 minutes on hold ≈ $4.40–7.30
Prior authorization worked end to end	$2.50	20–30 minutes ≈ $7.30–11.00
Denial appeal researched, packaged and filed	$4.00	30–45 minutes ≈ $11.00–16.50
Outreach conversation run by the agent	$0.12	—
Review reply drafted and published	$0.20	—
Campaign built and sent, per recipient	$0.04	—

8.3 The alternative: Revenue Recovery

Insurance-heavy practices that prefer pay-for-performance can choose 2.5% of collections attributable to MedOp's coding, denial-prevention, appeal and prior-authorization work, billed only against an auditable attribution ledger — no attribution, no invoice. MedOp's September 2026 total-cost review found comparable services priced at roughly

4–8% of collections at managed-RCM vendors and independent billing companies. Both models read the same work-unit ledger, so the choice costs MedOp little and is itself a fair-terms gesture.

8.4 Worked examples

Practice	Model	Monthly	Annual contract value	What they pay today (est.)
3-provider behavioral health, $60k/mo collections	Platform + AI Workforce meter	≈ $1,009	≈ $12.1k	≈ $1,574 (EHR + metered add-ons + processing)
Same practice	Platform + Revenue Recovery (2.5%)	≈ $2,067	≈ $24.8k	≈ $1,574
Single-location med spa, 2 injectors	Platform + meter; MedOp takes 0% of payments	≈ $378 + usage	≈ $4.5k + usage	Software ≈ $328 + outsourced Good Faith Exams + a payments markup on up to ≈ $117k/mo of revenue

Med-spa revenue uses AmSpa's 2023 average of $1.40M per location; the software line alone is close to the transparent category ceiling, so the aesthetics case rests on payments at cost and the compliance spine, not on the subscription. The 3-provider metered figure is the modelled example from MedOp's pricing decision record at typical volumes; at the heavier workload in Section 7.4 the same practice runs closer to $1,200. Incumbent cost estimates are from MedOp's September 2026 total-cost-of-ownership review of published price lists.

8.5 Why per-provider matters

Flat "unlimited providers" pricing caps contract value at the moment of sale and gives the largest discount to the practices most able to pay. Per-provider pricing with a metered labour line means MedOp's revenue grows as a practice grows and as it hands more work to the agents. That is the mechanism behind the net-dollar-retention figures the market now underwrites — Candid Health reported 180% net dollar retention with its July 2026 Series D, and Adonis above 130% with its March 2026 Series C — and MedOp is built to earn the same.

8.6 Unit economics

Variable cost is dominated by model inference and scales linearly with practice count; there is no per-tenant infrastructure beyond row counts. MedOp's cost model, measured at its heaviest realistic usage profile of roughly 1,000 agent calls per practice per day, puts gross AI cost near $745 a month per practice before the cost–quality cascade and $545–625 after it, with fixed infrastructure around $105 a month. That is a ceiling; typical practices run well below it. Three levers keep it there: prompt caching with 70%+ observed hit rates on practice context, fast-tier-first cascading for the fifteen non-safety agents, and a per-practice daily budget that degrades and then stops. The metered units are priced on the labour they replace, not on tokens: a denial appeal bills at $4.00 and costs roughly $0.21 in inference even on the most expensive model with reflection; a prior authorization bills at $2.50 against a few cents. Gross margin on the variable line is structurally high, and the platform fee carries the fixed cost.

8.7 Fair Terms Charter

Clause	Commitment
01	Month to month, always. No minimum term, auto-renewal, notice window or early-termination fee; cancel in the app.

Clause	Commitment
02	A warranted free one-click full export, complete within one business day, with a service credit if missed.
03	Zero price escalator for the life of a subscription, with published price history.
04	Everything in the number. No metered fee for any part of the infrastructure.
05	Payments at cost, shown as a formula. 0% platform fee, permanently.
06	No gag clause on reviews or benchmarking.
07	No licence taken over practice or patient data.

The charter is contract language, not marketing. Clause 02 ships as working software, load-tested at the largest tenant's volume, before it is advertised anywhere. A switching offer replaces the usual discount: an exit-fee credit of up to $1,500 against a documented termination or export invoice from the prior vendor, and a free parallel run until the practice cancels the incumbent.

9. Security, compliance and trust

MedOp is built to handle protected health information for behavioral-health practices, which are among the most sensitive records in medicine. Its controls are engineered into the platform and enforced by the database and the build pipeline, not by policy documents alone. This section states what is in place and what is still ahead.

9.1 Engineered controls

Tenant isolation at three layers

An application contract that requires a practice identifier on every query, audited in CI; Postgres Row-Level Security forced on 71 tables under a dedicated non-bypass role, so a query that forgets its filter returns nothing rather than another practice's data; and tenant-keyed cache invariants enforced by test.

Tamper-evident PHI audit trail

Every PHI read and write is logged by middleware; each row is hashed, sealed daily into a chain, verified by a scheduled job, and protected by an append-only trigger with a six-year retention floor (HIPAA §164.312(b)).

Identity and session

TOTP multi-factor login with hashed recovery codes; role-based access validated at every write; a 15-minute idle timeout; three principal classes — staff, patient, API client — with disjoint signing secrets that reject one another's tokens.

Encryption and secrets

TLS in transit; platform encryption at rest; nightly AES-256-GCM envelope-encrypted backups with a monthly restore drill; boot-time validation of required secrets with no environment fallbacks; secret scanning and CodeQL in CI.

42 CFR Part 2 and TCPA

Consent-gated disclosure of substance-use records with disclosure accounting; opt-out and 8am–9pm quiet hours enforced at the send chokepoint; a PHI-content guard on SMS and voice.

AI governance

Draft-only by default; confidence-band escalation as an opt-in; per-agent kill switches; a global draft-only switch; per-practice cost attribution enforced in CI; no patient-facing agents.

9.2 Operational resilience

Forty-five scheduled jobs run under a secret-protected prefix, each emitting a heartbeat watched by a dead-man's-switch monitor. Deployments run database migrations before promotion and verify the running tag afterwards. An incident-response runbook implements HIPAA §164.308(a)(6) with severity classes, a four-factor breach assessment and 60-day notification deadlines; a coordinated-disclosure policy commits to a one-business-day acknowledgement.

9.3 Status and roadmap — stated plainly

> **Where MedOp stands on attestation**
>
> MedOp's HIPAA technical safeguards are engineered, documented and self-assessed, and a SOC 2 policy and procedure pack with a Trust Services Criteria matrix has been drafted. **No third-party attestation — SOC 2, HITRUST or independent HIPAA assessment — has yet been obtained**, and the vendor Business Associate Agreement chain is institutionalised as a go-live gate that is completed per vendor before a practice is onboarded. A SOC 2 Type I engagement and an independent penetration test are the first items on the trust roadmap, alongside published per-agent evaluation scorecards and a public status page with historical uptime. The company would rather state this than imply otherwise.

10. Engineering: a platform you can verify

MedOp is built in the open with itself: every release is tagged, every material decision is recorded, and every claim in this section can be checked against the repository in minutes. The numbers below were measured on build v7.63.0 on 2 September 2026.

1,138 commits on main since 1 April 2026	**546** tagged releases, one changelog entry each	**45** architecture decision records	**21** CI gate workflows on every change
112 Prisma models · 62 migrations	**54** backend feature modules	**~186k** lines of TypeScript (backend + frontend)	**3,700+** automated tests, 80/80/80/70 coverage gate

10.1 Stack

Layer	Technology
Frontend	React 18, Vite, TypeScript 5, Tailwind, TanStack Query; installable PWA shell
Backend	Node.js, Express, TypeScript 5, Prisma; 54 feature modules; 23 repository modules that take the practice identifier as a required first argument
Database	Postgres on Neon with pgvector; Row-Level Security policies applied as versioned SQL

Layer	Technology
AI	Frontier large language models in three capability tiers (fast, standard, frontier), called through a single cost-attributed chokepoint under a Business Associate Agreement; a separate embedding model for retrieval; every model call logged with tokens, cost and outcome
Hosting and delivery	Vercel (api.medop.io on tag, app.medop.io on merge); 45 scheduled jobs; Resend and SMTP email lanes; Twilio SMS and voice; Sentry with a 56-field PHI scrubber
Integration	Public REST API v1 with OpenAPI; MCP server; HMAC-signed webhooks; FHIR R4 read import; one-way migration adapters

10.2 The gate suite

Twenty-one workflows run on every change. Beyond type-checking and the hermetic test suite, they include a tenant-scoping auditor that fails on any new cross-tenant query, a permission-coverage gate that fails on any unguarded route, a cost-attribution gate that fails on any unattributed model call, a migration-drift gate, an AI evaluation regression gate, secret scanning, CodeQL, Lighthouse and Playwright baselines, and an integration tier that runs the real RLS policies against a live Postgres and an ephemeral Neon branch. Twenty-one audit-style tests pin the properties that matter most: RLS policy coverage, no raw SQL on covered tables, no runtime DDL, registry–catalog consistency, cache tenant-safety and OpenAPI drift.

10.3 Decisions on the record

Forty-five architecture decision records document the load-bearing choices, from the system-of-record direction (ADR-0001) and the tenant-safety contract (ADR-0009) through the agent kill-switch architecture (ADR-0014), closed-loop agent improvement (ADR-0017), tamper-evident audit (ADR-0022), Part 2 segmentation (ADR-0026), confidence-band escalation (ADR-0030), database-enforced tenancy (ADR-0031), the patient portal and public API (ADR-0032, 0035), patient payments on Stripe Connect (ADR-0040) and the vertical-first go-to-market and pricing model this paper describes (ADR-0042). A buyer, partner or auditor reads the reasoning, not a summary of it.

11. Roadmap

The build plan is sequenced so that the market-facing difference is demonstrable before full depth lands, and so that nothing is marketed before it works.

Horizon	Milestone	What ships
Q4 2026	The Proof Scope (≈10 weeks)	Composable verticals on one record; the AI receptionist wired to the real availability engine; the behavioral-health minute-band rule; the aesthetics Good Faith Exam gate; cross-line containment proven by test; the warranted export; published pricing
Q4 2026	Commercial foundations	Provider seat bands in billing; the AI Workforce meter and provider-visible avoided-cost view; the Revenue Recovery attribution ledger; switching credit and self-serve signup; absorbed-bundle cost telemetry
Q1 2027	Behavioral health depth	Collaborative Care registry and time log with mutual-exclusivity enforcement; PCL-5 and licensed C-SSRS; telehealth modality and modifier logic; disclosure accounting export

Horizon	Milestone	What ships
Q1 2027	Aesthetics depth	Supervision attestation with state-rule evaluation; versioned consent and clinical photos; lot, expiry and witnessed wastage; memberships, packages, gift cards, commission and deferred revenue
Q2 2027	Therapy depth	The 8-minute engine and KX ledger; plan-of-care certification and progress-report gating; functional outcome instruments; therapy MIPS measures
Q2 2027	Trust and scale	Published per-agent evaluation scorecards; public status page; SOC 2 Type I; migration adapters for SimplePractice, TherapyNotes, WebPT, Jane, Aesthetic Record and Boulevard; founding-cohort self-serve readiness

Each vertical opens for general sale only when its sprints exit. Founding practices begin paying at a “production ready for your specialty” milestone rather than on a calendar date, so nobody pays for a demonstration.

12. The company

MedOp was founded by **Rachael Pereira** and **Dr. Paul A. Pereira**. **Rachael Pereira** is Chief Executive Officer and leads the company’s commercial direction, founding-practice programme and operations. **Paul Pereira**, Lead Development Engineer, brings more than twenty-five years across enterprise AI, healthcare SaaS, cybersecurity and aerospace, including engagements with NASA and the U.S. Space Force; he is the author of the platform’s codebase and of the architecture decision records that govern it. **Dr. Paul Alexander** serves as Security Officer, with responsibility for the HIPAA security programme, the SOC 2 readiness track, incident response and the vendor Business Associate Agreement chain described in Section 9. MedOp is headquartered in Florida, where its pilot practice — an integrated behavioral-health and medical-spa practice on the Treasure Coast — is the reference customer for the composition thesis this paper describes.

Role	Name	Responsibility
Chief Executive Officer	Rachael Pereira	Company direction, commercial model, founding cohort, operations
Co-Founders	Rachael Pereira · Dr. Paul A. Pereira	Founding thesis: one patient record across every line a practice runs
Lead Development Engineer	Paul Pereira	Platform architecture, the agent substrate, release engineering and the ADR record
Security Officer	Dr. Paul Alexander	HIPAA security programme, SOC 2 readiness, incident response, BAA chain

The company’s operating principle is written into its engineering standards: state what is true, cite it, and say “not yet” where that is the answer. This paper follows the same rule.

> **Talk to us**
>
> Practices: start at medop.io/pricing — published prices, a hands-on sandbox and a 30-day trial. Investors and partners: contact Rachael Pereira, CEO, via medop.io/demo or info@medop.io.

Appendix A — Agent catalog

The 27 operator-facing agents, their pod and the model capability tier they run on by default (fast, standard or frontier). Every agent runs on the shared substrate described in Section 5 and reports calibrated confidence on each action.

Agent	Pod	Tier	Job
Ambient Scribe	Clinical	Frontier	SOAP note generation from the encounter with E/M suggestion
Note Intelligence	Clinical	Standard	Summarise notes, extract action items, draft patient instructions
Care Gap Hunter	Clinical	Standard	HEDIS and MIPS measure tracking with closure actions
Risk Scoring	Clinical	Standard	Chronic-burden score, HCC suspects, readmission risk
eRx Safety	Clinical	Standard	Drug–drug interactions, allergy alerts, dosing
Compliance Monitor	Clinical	Frontier	Real-time HIPAA anomaly detection on PHI access
Compliance Audit	Clinical	Standard	Documentation adequacy against billed codes; OIG checks
Quality / MIPS	Clinical	Standard	MIPS composite score and improvement plan
Referral Manager	Front Office	Standard	Complete specialist referral packages
Staff Workflow	Front Office	Standard	Note-completion rates; burnout risk signals
Charge Capture	Revenue	Standard	Proposes E/M and CPT line items for signature
Grounded Coding	Revenue	Standard	CPT, E/M level and modifiers; underpayment risk
Denial Prevention	Revenue	Standard	Pre-submission denial risk and payer-specific fixes
Prior Auth	Revenue	Standard	Determines requirement; drafts clinical justification
Revenue Cycle	Revenue	Standard	Claim lifecycle triage; recovery estimates
Denial Appeal	Revenue	Frontier	Drafts the payer appeal letter for human signature
Underpayment Recovery	Revenue	Standard	Drafts dispute letters with contract-rate arithmetic
Financial Forecasting	Revenue	Standard	14- and 90-day cash-flow forecasts
Smart Scheduling	Front Office	Standard	No-show risk, demand forecast, reactivation slot ranking
Intake Intelligence	Front Office	Standard	Pre-visit demographics and insurance verification
Review Replies	Engagement	Frontier	HIPAA-safe public replies to patient reviews
Review Insights	Engagement	Standard	Sentiment clustering across reviews; alerts
Patient Outreach	Engagement	Standard	TCPA-compliant reminder and reactivation sequences
Campaign Intelligence	Engagement	Standard	Multi-segment campaigns with A/B arms and ROI
Patient LTV	Engagement	Standard	Lifetime value, churn risk, segmentation
Recall Engine	Engagement	Standard	Daily recall worklist with talking points

Agent	Pod	Tier	Job
Orchestrator	Orchestration	Standard	Plans multi-agent workflows as directed graphs; hands each step to the right specialist

Model tiers as configured at v7.63.0; the cost–quality cascade may run a lower tier first and escalate on calibrated confidence. The 37-type registry also carries the front-line dispatcher, internal document-triage and refill-queue automations, and the reflection and feedback meta-agents, which are not listed individually here.

Appendix B — Sources

Platform figures: MedOp repository at tag v7.63.0 (2 September 2026), CHANGELOG.md, docs/adr/ (45 records), docs/dd/, docs/plans/. Third-party figures:

AMA, “Physician burnout rate continues to decline, falling to nearly 42%,” Organizational Biopsy, 16 Apr 2026.
AMA, “Burnout way down, ‘pajama time’ stands still,” 13 Aug 2024.
Rotenstein et al., JAMA Network Open, via AMA, “Primary care visits run a half hour; time in EHR, 36 minutes,” 23 Jan 2024.
AMA, 2026 Physician Survey on Augmented Intelligence, 12 Mar 2026.
AMA, 2025 Prior Authorization Physician Survey (n=1,000), released 13 May 2026.
American Psychological Association, 2024 Practitioner Pulse Survey.
Ensora Health, “Practices Pulling Ahead” rehab-therapist survey (n>500), Apr 2026; 2026 Practice Success Report (n=1,124), 6 Aug 2026.
Kodiak Solutions, initial denial rate analysis (2,100+ hospitals, 300,000 physicians), 21 May 2025.
Premier Inc., “Claims adjudication costs providers $25.7 billion,” 24 Feb 2025.
Experian Health, 3rd annual State of Claims survey, 22 Sep 2025.
KFF, “Claims denials and appeals in ACA marketplace plans in 2024,” 27 Mar 2026.
KLAS Arch Collaborative, “Clinician EHR Experience 2025,” Jun 2025.
SAMHSA, 2024 National Survey on Drug Use and Health, 28 Jul 2025.
HRSA, State of the Behavioral Health Workforce 2025, Dec 2025.
IBISWorld, NAICS 621330 Offices of Mental Health Practitioners and NAICS 621340 Offices of PT/OT/SLP and Audiologists, 2025.
BLS, Occupational Employment and Wage Statistics, NAICS 621340 (May 2023) and national occupational medians (May 2025).
APTA, Benchmark Report: Hiring Challenges in Outpatient PT Practices, 16 Oct 2024.
Lewis et al., “Implementing measurement-based care in behavioral health,” JAMA Psychiatry, 2019.
Milliman for Meadows Mental Health Policy Institute, Collaborative Care uptake report, 13 May 2025.
American Med Spa Association, 2024 Medical Spa State of the Industry Executive Report recap, 6 Nov 2024.
Baird, “Takeaways from The Medical Spa Show 2025,” 20 May 2025.
Grand View Research, US Practice Management Systems Market (Jun 2025); US Behavioral Health Care Software & Services Market (Nov 2024).
Menlo Ventures, “2025: The State of AI in Healthcare,” 21 Oct 2025.
Rock Health, 2025 year-end digital health funding overview, 12 Jan 2026.
athenahealth, native ambient documentation general-availability announcement, 13 Aug 2026.
Candid Health, $120M Series D announcement, 21 Jul 2026; Adonis, $40M Series C announcement, 25 Mar 2026.
CMS Claims Processing Manual Ch. 5 §20.2; CY2026 therapy threshold; 42 CFR Part 2 Final Rule (89 FR 12472); ASTP/ONC certification programme; HTI-5 proposed rule (90 FR 60970); 2026 MIPS eligibility thresholds.
Published price lists and terms of SimplePractice, TherapyNotes, Jane, WebPT, Zenoti, Boulevard, Aesthetic Record and Mangomint, reviewed Sep 2026; Helcim interchange-plus benchmarks, Feb 2026.

This white paper describes MedOp’s platform as of build v7.63.0 and its published commercial model as of September 2026. Capabilities described as planned, scheduled or in the Proof Scope are not yet generally available. Illustrative value figures are models built from cited public benchmarks, not measured customer results. Nothing in this document is an offer of securities.