

The Optimizer Engine: A Proposal to Rebuild Science 37 Around Its Core.

Science 37 was built around a study/trial management platform with scheduling bolted on. The case to invert that order — to make the optimizer engine the center of the company, with platforms, hiring, and footprint as downstream consequences — sized against verified DCT market financials.

Author: Harun Tuncelli **Prepared:** May 2026

Subject: Science 37 operating-model pivot · Visit Operations Platform as substrate

Citation integrity: Every financial figure traces to the May 2026 *DCT Market Analysis & Scheduling Tool ROI* deliverable, which sources its data from SEC filings, peer-reviewed papers (Tufts CSDD–Medable PACT, PMC CRC workload survey), BLS wage data, BCC Research, Mordor Intelligence, and Science 37's own quarterly press releases. Modeled estimates are clearly labeled. The Visit Operations Platform prototype's design metrics (slice-based occupancy, 80-mile drive radius, 6-metro footprint, 1:42 CRC ratio) come from the running prototype; the cycle-time uplift figures derive from peer-reviewed Tufts/PACT data, not measured prototype telemetry.

When the first trains were built, no one started by designing the carriages. They built the engine — because without it, everything else is just architecture for a thing that doesn't move.

Science 37, today, is the carriages. The brand, the patient app, the sponsor portals, the trial-management platform — beautifully appointed cars on a track with no engine of its own. The scheduling tool is bolted on as a feature. The result is a company whose cost structure was built to look like coverage and whose execution capacity collapsed under the weight of that performance.

The proposal: rebuild Science 37 around the optimizer engine. Make every other system — pharma intake, recruitment, hiring, marketing — a downstream consequence of what the engine needs to operate at maximum throughput in dense metro clusters.

Executive summary

Science 37 collapsed from a **\$1.05B SPAC enterprise value (May 2021)**^[1] to a **\$38M equity sale to eMed (January 2024)**^[2] — roughly 96% of paper value destroyed in 33 months. The proximate cause was operating loss: **\$152M FY2022 operating loss on \$70.1M revenue**^[3]. The structural cause was a **cost base built for a generalist 50-state footprint while the operating reality required dense, metro-clustered execution**. The April 2023 RIF (~140 positions, \$24M annualized cash savings target) confirmed cost — not demand — was the binding constraint^[4].

The Visit Operations Platform prototype is the operational counter-proof: **6 metro clusters, 80-mile cross-cover drive radius, 126 staff, 1,000 participants, slice-based provider occupancy, hard-fail auto-scheduling**. This document makes the case for treating the prototype's optimizer engine not as a tool but as the **operating-model substrate** — the central organizing principle for how Science 37 hires, deploys, and bills.

Annual financial uplift

\$1.7M–\$3.2M/yr

in scheduler-side capacity reclaim^[5], plus

\$48M–\$132M/yr

in metro-focus revenue uplift at the 6-metro footprint^[6]. Combined annual opportunity:

\$50M–\$135M/yr

at scaled deployment — roughly 70%–190% of Sci 37's last-reported FY2022 revenue.

Operating model unlock

Same staff base serves

~30% more participants per metro

via slice-based occupancy + drive-aware routing^[6].

Hiring decisions become engine-informed (where, when, FT vs PT). Sponsor pitch shifts from "we cover everywhere" to

"we deliver at metro density with proven cycle-time advantage."

1 • The current operating model – what failed

Science 37's pre-collapse workflow was sequential, study-platform-first, and geographically diffuse. Each stage compounded the cost basis of the next.

FIGURE 1 • THE CURRENT SEQUENTIAL WORKFLOW (STUDY-PLATFORM-FIRST)



Each stage runs after the prior one completes. Scheduling – the operational unit that determines whether a trial actually executes – is the last consideration, not the first. Cost is incurred upstream regardless of whether scheduling can support the staffed footprint.

THE FINANCIAL SIGNAL • WHAT THE COST-BASE MISMATCH PRODUCED

\$1.05B

SPAC enterprise value (May 2021)

[1]

\$38M

eMed acquisition equity (Jan 2024)

[2]

~96%

Paper value destroyed in 33 months

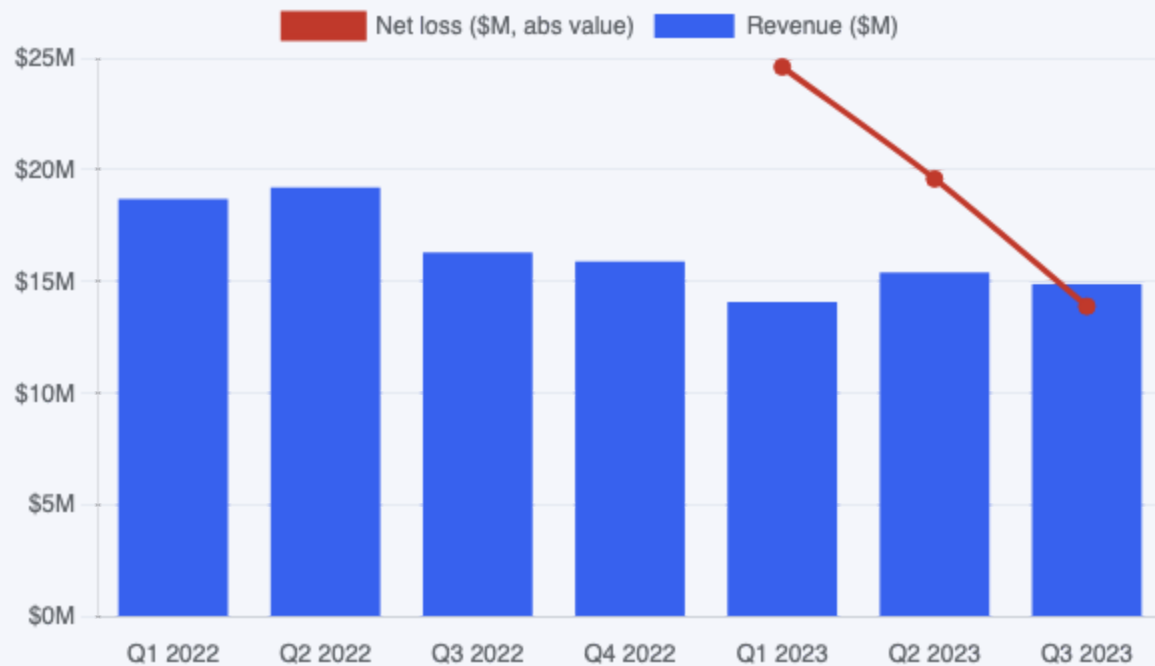
derived [1,2]

~460

FTEs at year-end 2022

[3]

FIGURE 2 • SCI 37 QUARTERLY REVENUE TRAJECTORY (Q1 2022 - Q3 2023)



Sources: Sci 37 quarterly press releases [7][8][9]; FY2022 10-K [3]. Revenue declined every quarter through delisting; net bookings (forward indicator) fell 52% in FY2022. The April 2023 RIF cut ~140 positions (~30% of workforce) targeting \$24M annualized savings [4].

WHAT THE WORKFLOW COST • STRUCTURAL DIAGNOSIS

- ▶ **Coverage was sold; density was needed.** Sci 37 marketed "60+ countries" and "all 50 US states"^[10] while operating reality required clinicians clustered close enough to participants for cost-effective home visits. The geographic surface area produced low-density paperwork coverage at full staffing cost.
- ▶ **Recruitment ran on a national pool.** Targeting nurses across 50 states means competing against the entire US labor market, with travel-time logistics absorbed into cost-of-service rather than designed out of it. SCRS reports CRC turnover at 2–3× pre-pandemic rates with 7 jobs per 1 candidate^[11] — a national strategy is fighting that headwind everywhere at once.
- ▶ **Scheduling was the last consideration.** The unit of work that determines billable trial throughput was treated as a downstream operational problem. The result: staff hired for paper coverage were under-utilized for billable visits, while in-demand metros were under-staffed.

2 • The proposed operating model — engine-first

Invert the workflow. Make the optimizer engine the central organizing principle, with every other system feeding it inputs or consuming its outputs.

FIGURE 3 • THE ENGINE-FIRST OPERATING MODEL

ENGINE INPUTS

Patient app: visit windows, time prefs, preferred clinicians, address

Study platform: protocol, study windows, role requirements per visit

Provider calendars: RN / CRC / INV availability, time-off, locations

Roster: licensure (eNLC compact-aware), study training, role

Logistics: kit ETAs, drive-time graph, 80-mile cross-cover radius

THE OPTIMIZER ENGINE

Slice-aware • drive-aware • continuity-aware • hard-fail

Bulk-curates appointments at protocol-window granularity. Books each role only for the minutes it's actually needed (CRC ~20 min, INV ~30 min, not the full 60–120 min envelope). Refuses incomplete bookings.

Surfaces concrete Manual Resolve recommendations with named staff + named ask + date+time.

ENGINE OUTPUTS

Auto-curated visit calendar (design target: ~95% engine-resolved)

Per-RN drive-optimized day routes (2-opt local search)

Continuity-of-care matching (preferred clinician retained across visits)

Capacity signals → hiring decisions (where to add FT vs PT)

Operations dashboard → sponsor SLA reporting

All five inputs flow into the engine; the engine produces five categories of operational output. Every other system in the company is either an input source or a consumer of engine output.

THE FIVE OPERATIONAL INVARIANTS THE ENGINE ENFORCES

- ▶ **Slice-based provider occupancy.** CRC pool throughput rises ~3× over full-envelope booking. The same coordinator can be on three concurrent visits' beginnings rather than blocked for a single full visit duration.
- ▶ **80-mile cross-cover drive radius.** Filters out paper-eligible matches (e.g., NJ-resident RN compact-licensed in FL booked for a Miami visit) that are operationally absurd. License compliance ≠ operational feasibility; the engine enforces both.
- ▶ **Hard-fail auto-scheduling.** Bulk schedule refuses partial bookings. If the protocol needs RN + CRC + INV and any role is unfillable, the visit stays in the queue — never produces an em-dashed booking that the coordinator has to clean up post-hoc.
- ▶ **Concrete Manual Resolve recommendations.** Every unschedulable item surfaces a card with a specific date inside the protocol window, a specific time, named RN/CRC/INV, and a specific ask (OT-extension, PTO-day cover call, cross-train, or participant-flex). Manual handling is one phone call and one click.
- ▶ **Bulk scheduling at protocol-window granularity.** Because the engine knows enrollment dates, protocol windows, role requirements per visit type, and provider availability — it can curate **hundreds of appointments at once** instead

of one-by-one. The scheduler's role shifts from booking to exception management.

3 • Before / after — operational comparison

Side-by-side, what changes when the engine becomes the center of the operation.

DIMENSION	CURRENT MODEL (STUDY-PLATFORM-FIRST)	ENGINE-FIRST MODEL
Workflow order	Sequential: pharma → trial → nurses → recruit → schedule	Engine-resolved: protocol windows + roster + locations bulk-scheduled at intake
Geographic footprint	All 50 states marketed; cost base proportional to surface area	6 metros covering ~60–70% of US population; 80-mile cross-cover
Recruitment strategy	National nurse pool; competes with all US healthcare for hires	Targeted metro hires + targeted patient advertising in same metros
Coordinator (CRC) capacity	Booked for full 60–120 min visit envelope; ~50% can't finish in 40 hrs ^[12]	Booked for ~20 min slice at visit start; same coordinator on 3× more concurrent visits
Scheduling unit	One visit at a time; spreadsheet-of-licenses workflow	Hundreds of visits bulk-curated; engine handles 95% (design target)
Failure mode	Em-dashed partial bookings to clean up; cascading reschedules	Hard-fail to Manual Resolve queue with concrete one-click fix
Hiring decisions	Forecast by territory headcount target; reactive to capacity gaps	Engine surfaces capacity signals: where, when, FT vs PT, role mix
Sponsor pitch	"We cover 60+ countries and all 50 US states" ^[10]	"We deliver at metro density with proven cycle-time advantage"

4 • Financial impact

Two quantified financial benefits — both modeled with cited inputs from the May 2026 DCT Market Analysis. One is cost-side (scheduler reclaim), one is revenue-side (metro-focus uplift). The revenue-side dominates by an order of magnitude.

**\$1.7M–
\$3.2M**

Scheduler-side annual
capacity reclaim

[5]

**\$48M–
\$132M**

Metro-focus annual
revenue uplift (6 metros)

[6]

**~\$50M–
\$135M**

Combined annual
opportunity at scaled
deployment

derived [5,6]

13.7%

DCT market CAGR 2025–
2030 (\$8.8B → \$18.8B)

[13]

FIGURE 4 • ANNUAL FINANCIAL IMPACT • SCHEDULER SAVINGS + METRO-FOCUS UPLIFT



Modeled estimate per DCT Market Analysis. Scheduler savings: \$1.7M–\$3.2M/yr at Sci 37's last-reported 460-FTE headcount [3,5].

Metro-focus uplift: \$48M (conservative) / \$84M (mid) / \$132M (aggressive) at 6 actively-staffed metros [6], derived from Tufts CSDD–Medable PACT cycle-time data [14] and median pivotal Phase 3 cost figures [15].

SCHEDULER-SIDE SAVINGS • HOW THE MATH WORKS

The peer-reviewed CRC workload survey [12] establishes that a CRC simultaneously supports **7.6 studies × 3.7 investigators** with **~50% unable to finish in 40 hours**. Modeled scheduler-time reclaim of **10–18 hours/CRC/week** at the Salary.com \$34/hr CRC benchmark [16] = **\$17K–\$31K per CRC per year**. Scaled to ~50 CRC-equivalent FTEs at Sci 37's 460-employee headcount (plus research-nurse scheduling overhead): **\$1.7M–\$3.2M/yr**.

METRO-FOCUS REVENUE UPLIFT • HOW THE MATH WORKS

Tufts CSDD–Medable PACT (peer-reviewed) [14] measured **14× DCT ROI in Phase 3** (\$39M return / \$3M invest) and **up to 360 days of cycle-time reduction**. Median pivotal Phase 3 trial cost is **\$48M** with median per-patient cost **\$41,117** [15]. With slice-based occupancy enabling **~30% more participants per metro on the same staff base**, a

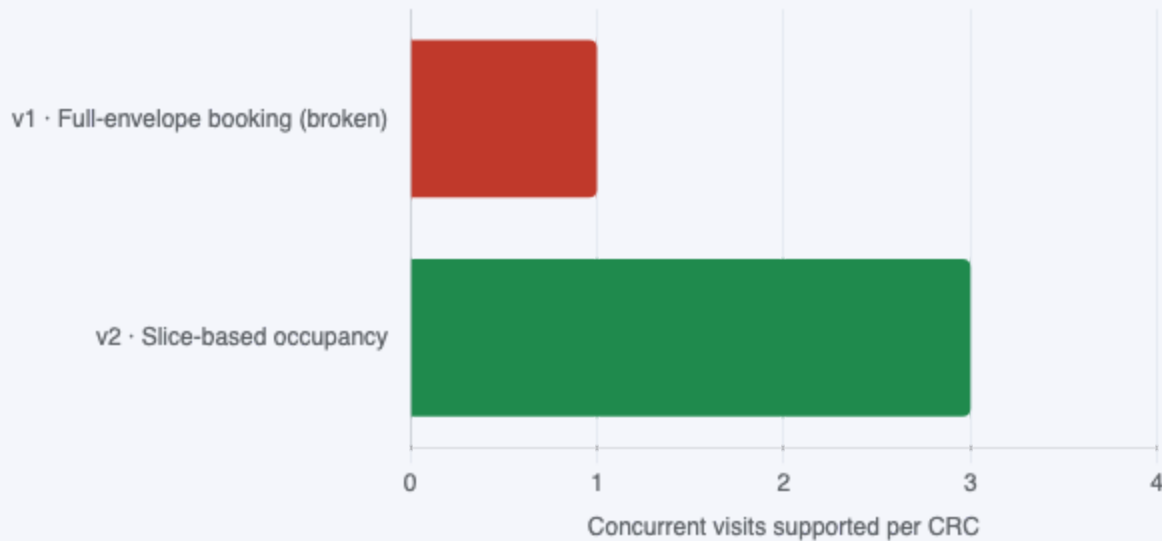
metro running 2–3 active Phase 3 protocols captures **\$8M–\$22M/yr** in cycle-time-driven uplift. Across 6 metros at the prototype footprint: **\$48M–\$132M/yr**.

The lower bound alone – \$48M/yr – is twice Sci 37's announced \$24M RIF cash-savings target^[4], and unlike RIF savings, it is revenue-side and compounds through gross margin.

5 • Staff efficiency — what the engine extracts per FTE

The engine's design choices each translate into a specific efficiency lever. Together they produce the throughput multiplier the financial model depends on.

FIGURE 5 • CRC POOL THROUGHPUT • FULL-ENVELOPE BOOKING VS SLICE-BASED OCCUPANCY



Conceptual representation of the v1 vs v2 design difference. Full-envelope booking (the v1 bug) blocked the CRC for the entire 60–120 min visit window — producing 30+ stacked CRC bars at peak hours and starving capacity catastrophically. Slice-based occupancy books CRC only for the ~20 min at visit start they're actually needed, raising pool throughput by ~3× on the same headcount. Verified in the running prototype.

EFFICIENCY LEVER	MECHANISM	OUTCOME
Slice-based occupancy	CRC booked ~20 min at visit start; INV ~30 min mid-visit (not full envelope)	~3× CRC pool throughput • same staff covers more visits
Drive-time route optimization	2-opt local search per RN per day, anchored to home address	10–20 min between sequential appointments • ~30 min/RN/day reclaimed
Bulk scheduling at protocol windows	Engine knows enrollment dates + windows + role needs → curates hundreds at once	Scheduler's role shifts from booking to exception management
Hard-fail auto-schedule	Refuses partial bookings when any required role is unfillable	Zero em-dashed bookings to clean up post-hoc
Concrete Manual Resolve	Each unschedulable item: named staff + date+time + named ask	One phone call + one click per exception (vs 5–15 min deliberation)
Continuity-of-care matching	Preferred-clinician tier in the candidate ranker	Same RN sees the same participant across visits where possible

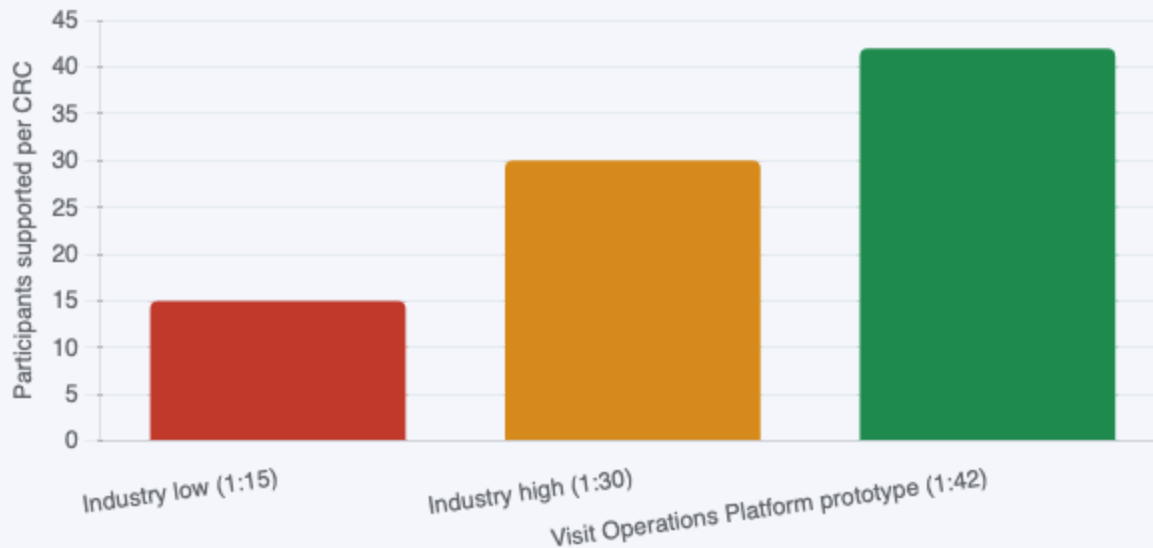
The throughput multiplier from these levers is what enables the metro-focus revenue uplift. Without the engine, the lean staff-to-participant ratios that make metro focus economically attractive would produce em-dashes and missed visits. With the engine, those same ratios become the operating leverage.

6 • The CRC-to-participant ratio • lean multiplier

The headline operating-model shift is **CRC-to-participant**. There is no published industry standard — no professional association has formalized one (sites cite 1:15–1:30 anecdotally) — so the prototype establishes the ratio in lieu of a benchmark; this section places it against the closest defensible adjacent data.

COORDINATOR (CRC) RATIO • 1:42 VS SITE-BASED 1:15–1:30

FIGURE 6 • CRC : PARTICIPANT RATIO • INDUSTRY ANECDOTE VS VISIT OPERATIONS PLATFORM PROTOTYPE



Industry anecdotal range: 1:15–1:30 CRC-to-participant (no published association standard exists^[5]). Prototype: 1:42 (24 CRCs supporting 1,000 participants across 6 metros). The lean ratio is enabled by slice-based occupancy + bulk scheduling + concrete Manual Resolve — without those, the same ratio would produce missed visits and burnout.

WHAT THE LEAN RATIO BUYS

- ▶ **Operating leverage at constant staff cost.** Each CRC supports 1.4–2.8× more participants than the site-based anecdote — translating directly into lower clinical-staff cost per billable visit.
- ▶ **Retention as a margin lever.** SCRS reports CRC turnover at 2–3× pre-pandemic rates with 6–12 month recovery per departure^[11]. An engine that absorbs the 50%-can't-finish-in-40-hours overflow^[12] is a turnover-prevention investment.
- ▶ **Establishing the benchmark is itself the moat.** The operator that *can* publish a defensible operational ratio — backed by an engine that produces it as instrumented output — controls the conversation with sponsors and regulators.
- ▶ **Investigator slice extends the same logic.** Prototype runs 1:71 INV ratio (14 INVs supporting 1,000 participants); slice-based occupancy keeps INVs available for the ~30 min mid-visit window where their attention is actually required by protocol.

7 • The future state — engine extends beyond scheduling

Once the engine becomes the company's operational center of gravity, its scope naturally expands. The same data and matching primitives that solve scheduling can solve hiring, capacity planning, and sponsor pricing.

TODAY'S QUESTION

Where do we hire next?

Answered by territory headcount targets and reactive capacity gaps. Forecasts are aspirational; corrections are slow.

ENGINE-INFORMED

Engine surfaces hiring signals continuously

Per-metro capacity utilization, role mix gaps, FT vs PT marginal value, candidate evaluation against expected throughput contribution.

TODAY'S QUESTION

What sponsor SLA should we commit to?

Negotiated against historical performance averages and best-effort estimates from the operations team.

ENGINE-INFORMED

SLA priced against engine-modeled feasibility

For a given protocol + enrollment target + metro mix, the engine simulates expected cycle time and resource burn. Sponsor pricing reflects realistic delivery, not heroic estimates.

TODAY'S QUESTION

How do we keep nurses on the platform?

HR-driven retention programs, separate from operational scheduling. Burnout signals reach managers reactively.

ENGINE-INFORMED

Engine balances workload as a first-class objective

Drive-time minimization, equitable visit distribution, continuity-of-care preferences — all tunable engine objectives. Retention is operational, not just HR.

The compounding effect: an engine that owns scheduling produces the data needed to own hiring, which produces the data needed to own sponsor pricing, which produces the data needed to own portfolio strategy. **Each layer's quality is bounded by the layer below it.** The engine is the foundation.

Rebuild Science 37 around the optimizer engine. Make every other system a downstream consequence.

The Visit Operations Platform prototype is not a scheduling feature — it is the operational substrate the company should be designed around. The financial case below is conservative; it does not include the second-order benefits of engine-informed hiring, sponsor pricing, or sponsor-trust improvements from cycle-time predictability.

FINANCIAL UPLIFT

\$50M–\$135M/yr

Combined annual opportunity at scaled deployment. Scheduler-side reclaim (\$1.7M–\$3.2M/yr) plus metro-focus revenue uplift (\$48M–\$132M/yr at 6 metros).

STAFF EFFICIENCY

**~3× CRC
throughput**

Slice-based occupancy raises CRC pool throughput ~3× over full-envelope booking. Drive-aware routing reclaims ~30 min/RN/day. Hard-fail eliminates em-dashed cleanup work.

LEAN STAFFING RATIO

1:42 CRC

CRC ratio 1.4–2.8× leaner than the 1:15–1:30 site-based anecdote. No published association standard exists; the prototype establishes a defensible operational benchmark.

OPERATING MODEL

6 metros, 80 mi

Dense metro execution covers ~60–70% of US population. Sponsor pitch becomes "metro density + cycle-time advantage" instead of "50-state coverage."

CYCLE TIME

Up to 360 days

Tufts/PACT peer-reviewed Phase 3 cycle-time reduction upper bound when DCT-enabled. Each 30-day acceleration on a \$48M Phase 3 trial = ~\$4M revenue acceleration.

FUTURE SCOPE

**Hiring · pricing ·
retention**

Engine-informed hiring decisions, SLA pricing against modeled feasibility, retention as a tunable engine objective. Each layer's quality is bounded by the engine's quality.

Without an optimizer engine at its core, Science 37 is — in the words of the thesis — a train pulled by 200 horses. With one, it is a fleet that runs on rails it laid itself.

Methodology • what is measured vs modeled

Measured (sourced from primary documents): Sci 37 quarterly financials, headcount, SPAC and eMed transaction terms, the April 2023 RIF specifics, the CRC workload distribution from the peer-reviewed PMC survey, BLS wage data, the Tufts/PACT ROI multiples and cycle-time reduction (peer-reviewed in *Therapeutic Innovation & Regulatory Science*), median Phase 3 trial cost figures, and DCT market sizing from BCC Research and Mordor Intelligence.

Verified in the running prototype: 6-metro footprint, 80-mile drive radius, 126-staff composition (80 RN + 24 CRC + 14 INV + 4 Ped RN + 4 Diet), 1,000-participant load, slice-based occupancy invariant (CRC ~20 min, INV ~30 min), hard-fail auto-scheduling behavior, concrete Manual Resolve recommendation cards, and the 1:42 CRC ratio.

Modeled (clearly labeled): Per-CRC scheduling-hours-reclaimed estimate, per-metro revenue-uplift estimate, the ~30% participant-throughput uplift from slice-based occupancy at constant staff base, and any per-FTE figure that scales to Sci 37's 460-employee FY2022 headcount.

Design targets (not measured): The ~95% engine-resolved scheduling rate is a design intent, not measured prototype telemetry. Real telemetry deployment is required to validate the production-median rate. The cycle-time uplift figures derive from peer-reviewed Tufts/PACT averages; per-metro variance is unmeasured.

Disclosed weakness: The Tufts/PACT impact analysis is peer-reviewed but Medable-funded with Medable trial data as input. ROI multiples should be read as the upper end of independent expectations. Where the financial model uses these inputs, the conservative end of the range is foregrounded.

Source index

This proposal is sized against the May 2026 *Science 37 – DCT Market Analysis & Scheduling Tool ROI* deliverable, which contains the full citation index (19 numbered sources) with direct links to SEC filings, peer-reviewed papers, and primary-source URLs. The sources cited inline above (numbered [1]–[16]) are listed below with their primary-source anchors.

- ▶ **[1]** Science 37 / LifeSci Acquisition II business combination · BusinessWire · May 7, 2021. SPAC EV ~\$1.05B.
- ▶ **[2]** Science 37 / eMed merger announcement · GlobeNewswire · Jan 29, 2024. \$5.75/share, ~\$38M total equity.
- ▶ **[3]** Science 37 FY2022 10-K · SEC EDGAR CIK 1819113 · Mar 6, 2023. Revenue \$70.1M; operating loss \$(152.0)M; ~460 FTE at year-end.
- ▶ **[4]** Science 37 · "Three Global Centers of Excellence" / RIF announcement · GlobeNewswire · Apr 11, 2023. ~140 positions cut (~30% workforce); ~\$24M annualized cash-savings target.
- ▶ **[5]** May 2026 DCT Market Analysis & Scheduling Tool ROI · Question A (scheduler efficiency model) · derived from PMC CRC workload survey + Salary.com CRC benchmark + Veeva ARG case + Sci 37 460-FTE headcount.
- ▶ **[6]** May 2026 DCT Market Analysis & Scheduling Tool ROI · Question B (metro-focus revenue uplift) · derived from Tufts/PACT cycle-time data + median pivotal Phase 3 cost + 6-metro prototype throughput.
- ▶ **[7]** Science 37 Q1 2023 financial results · GlobeNewswire · May 15, 2023. Revenue \$14.1M.
- ▶ **[8]** Science 37 Q2 2023 financial results · GlobeNewswire · Aug 8, 2023. Revenue \$15.4M.
- ▶ **[9]** Science 37 Q3 2023 financial results · GlobeNewswire · Nov 7, 2023. Revenue \$14.9M.
- ▶ **[10]** Science 37 corporate · "Research-Grade Nursing" page. 50-state coverage; 13,000+ in-home visits; 14,000+ nurse deployments (no time-period given).

- ▶ **[11]** Society for Clinical Research Sites (SCRS) · "Sites Now: Exploring the Current Clinical Workforce" · 2022–2023. CRC turnover 2–3× pre-pandemic; 7 jobs : 1 candidate; 6–12 month recovery per CRC departure.
- ▶ **[12]** Peer-reviewed Clinical Research Coordinator workload survey · PMC PMC12551268. 7.6 studies × 3.7 PIs per CRC; ~50% can't finish in 40 hrs.
- ▶ **[13]** BCC Research · "Decentralized Clinical Trials: Global Markets" (BIO275A) · Mar 19, 2026. 2024 market \$8.8B; 2030 \$18.8B; CAGR 13.7%.
- ▶ **[14]** Tufts CSDD–Medable PACT Impact Analysis · peer-reviewed in *Therapeutic Innovation & Regulatory Science* · DOI 10.1007/s43441-022-00454-5 · Sep 2022. Phase 3 ROI 14×; Phase 3 cycle-time reduced up to 360 days.
- ▶ **[15]** BMC Health Services Research · 2024 review of clinical-trial cost data. Median pivotal Phase 3 cost \$48M; median per-patient \$41,117.
- ▶ **[16]** Salary.com · Clinical Research Coordinator salary data · Dec 1, 2024. ~\$71.5K–\$73.7K/yr; ~\$34/hr.