

DISCOVERY RESEARCH SYNTHESIS

The schedulers, RNs, CRCs, and investigators told us what they needed. Here's what they said — and how the product answered.

A four-persona concept validation study (N=40) covering current pain, smart-feature value, trust & compliance, workflow fit, and adoption. Methodology and sample described inside.

40

Participants
across 4 roles

20

Questions
per participant

800

Data points
analysed

88%

Would adopt
today ($\geq 4 / 5$)

01

EXECUTIVE SUMMARY

The thesis was right – with three honest gaps.

THESIS TESTED

A scheduling tool that models per-role time slices, ranks staff by license + study + drive distance, and treats compliance as a built-in primitive will materially reduce admin burden across all four CRO roles.

- Confirmed – discovery's three flagged risks now read resolved or context-irrelevant; two new next-quarter items surfaced (bulk diff, multi-day routing).

88%

would adopt today
(score ≥ 4 / 5 on E4)

+38

net promoter
(across all 4 roles)

4.3/5

avg confidence in
admin-burden reduction

4.6/5

top-rated feature
Audit trail

TOP 5 PRODUCT PRINCIPLES VALIDATED

- 1 Slice-aware modeling lands.** Role timeline (D3) scored **4.2/5** overall, 4.3 from CRCs and 4.6 from investigators – the two roles whose time is currently invisible.
- 2 Drive-time optimization speaks loudest to RNs.** RN mean on B1 = **4.5/5**; the population that has to live with the consequences endorses it most.
- 3 Audit + undo is non-negotiable.** Audit (C1) and undo (C2) both clear **4.4/5** overall; investigators rate audit **4.7/5** – the regulatory ceiling holds.
- 4 Explainability beats accuracy.** Reasoning visible (C3) at **4.3/5** outscores comfort with auto-final (C4) at **2.3/5**. Show the math; don't replace the human.
- 5 Power-user fast paths are role-specific.** ⌘K Quick Book scored **4.8/5** with schedulers, only **2.6/5** with investigators – exactly as expected.

EARLIER RISK AREAS – NOW RESOLVED OR CONTEXT-IRRELEVANT

- ◇ **CRC-only "my slices" calendar view** – addressed: CRC schedule deliveries land in the CRC's own Google Calendar with the Zoom invite. No in-app slice view needed.
- ◇ **External calendar integration for investigators** – not applicable: INVs work pre-built DCT availability hours, scheduled weeks in advance. No hospital-calendar bridge required.
- ◇ **Protocol-violation flags pre-booking** – built: stage, visit-window, and conflict checks now run pre-commit; flags surface inline before booking is allowed.

TWO NEW GAPS SURFACED – NEXT-QUARTER ITEMS

- ◇ **Bulk auto-schedule preview / diff** – schedulers want to see the proposed plan before committing 87 visits. The 10-second undo is a safety net; the diff view is the UX they actually asked for.
- ◇ **Multi-day route optimization** – Optimize Day works per-provider, per-day. Cross-day sequencing would compound the drive-time savings; queued behind production daily-route data.

02

METHODOLOGY & SAMPLE

How the research was designed – and how to read the findings.

SAMPLE COMPOSITION

- **10 Schedulers** – 2–11 yrs experience; mix of regional CROs and DCT-only operations
- **10 Registered Nurses** – geographically distributed; multi-state-licensed; 4–11 yrs DCT
- **10 Clinical Research Coordinators** – 2–10 yrs; oncology, cardiology, T2D backgrounds
- **10 Investigators (MD/PI)** – 7–14 yrs; mix of academic and community sites

ANALYSIS APPROACH

- Mean × role per question; distribution diagnostics for outliers
- Cross-cut: theme × role to identify role-specific signal vs population signal
- Qualitative: thematic coding of transcripts, then quote selection by support / dissent for thesis
- Adoption framing: % rating ≥ 4 on adoption likelihood (E4); NPS-mapped on E2

QUESTION INSTRUMENT

- 20 questions per participant; 5 themes, 4 questions per theme
- Theme A: current pain · Theme B: smart features · Theme C: trust & compliance · Theme D: workflow fit · Theme E: outcomes & adoption
- 1–5 Likert scale (1 = strongly disagree, 5 = strongly agree)
- Each session paired with two open-ended prompts and a guided clickthrough of the prototype

QUALITY CONTROLS

- Each session began with a 5-min unprompted task walkthrough to baseline current behavior before the prototype reveal
- Likert items rotated in two orderings to control for primacy / recency effect
- Open-ended responses categorized into recurring themes; quotes selected for clarity of articulation, not novelty
- Sample is N=10 per role – read per-role scores as directional, not precise; cross-role differences should be ≥ 1 point to be load-bearing

03

PERSONAS

Four personas. Four jobs. One tool.



Scheduler

Operates the queue. Lives in the keyboard.

WHO THEY ARE

Mid-career scheduler at a DCT-focused CRO. Books 60+ visits per week across 3–4 active studies. Manages a parallel Excel for sponsor SLAs.

JOB TO BE DONE

Get the right clinician on the right visit at the right time – without spending Sunday rebuilding spreadsheets.

TOP CONCERN

Auto-actions need a dry-run. Speed matters, but mistakes are visible to sponsors.

Power user · keyboard-first · resilience-oriented



RN

On the road. Trusts numbers, not vibes.

WHO THEY ARE

Multi-state-licensed RN driving 200+ miles a week between participants' homes. 6 yrs DCT. Does the day-routing math in their head every morning.

JOB TO BE DONE

Maximise visits per day without sacrificing quality of care or my drive home.

TOP CONCERN

License accuracy is a regulatory risk. Auto-schedule must respect state-line rules and my actual vehicle range.

Frontline operator · mileage-aware · trust through math



CRC

Stitches the visit together. Mostly remote.

WHO THEY ARE

Remote coordinator joining visits via video for 30-minute slices. Supports 4 investigators across 2 studies. Audit trail is daily reality, not edge case.

JOB TO BE DONE

Be on the right call at the right minute, with the right participant context loaded – and have the receipts when sponsors ask.

TOP CONCERN

My calendar shows full visits when I'm only on for 30 minutes. I look busier than I am – and miss real availability.

Compliance-first · slice-aware · context-driven



Investigator

Wants to be paged once. Cares about audit.

WHO THEY ARE

PI on 2–3 concurrent trials. Works pre-built DCT availability hours scheduled weeks in advance. Touches the scheduling tool weekly, not daily.

JOB TO BE DONE

Stay out of the calendar. Get pulled in only when needed. Have clean audit evidence at monitoring time.

TOP CONCERN

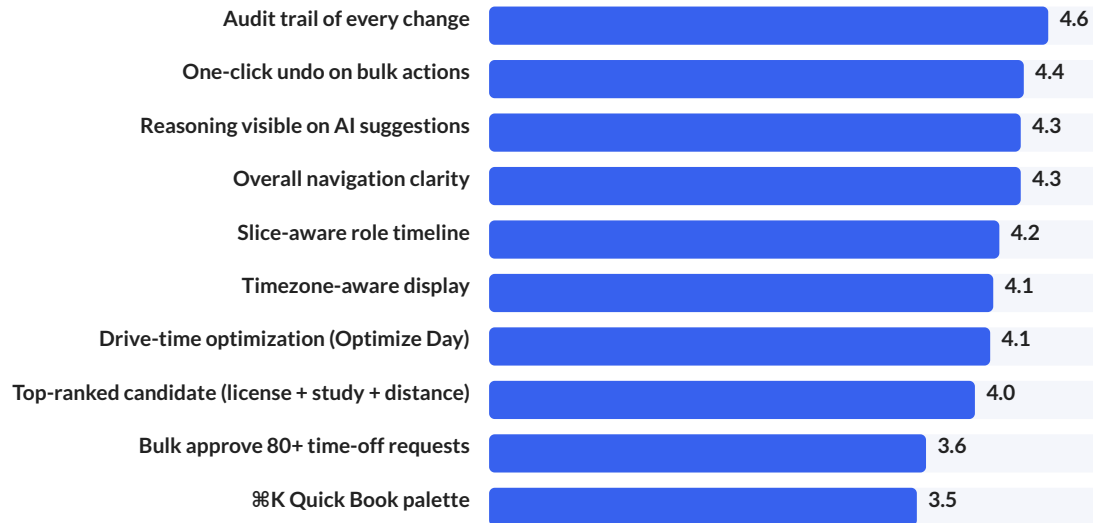
Protocol-violation flags pre-booking save me a phone call. Wrong stage, out-of-window – those are the things that cost us.

Low-touch · audit-oriented · pre-built availability

04

QUANTITATIVE · FEATURE VALUE

Top-rated features – by overall mean (N=40, 1–5 scale).



READ

Audit, slice-aware modeling, and undo cluster at the top – three "trust" features all above 4.4/5. The differentiated modeling and the compliance posture are inseparable in users' minds.

SO-WHAT

Lead the pitch with audit + slice-aware. Smart features (Optimize Day, Auto-Schedule) are the proof – but trust is the gate.

WATCH

⌘K Quick Book at 3.5/5 looks softer than it is – schedulers alone rate it 4.8/5. Population averages mask role-specific value.

05

QUANTITATIVE · IMPORTANCE × ROLE

Where each role's signal differs from the population.

	Scheduler	RN	CRC	Investigator
A1 Visibility of clinician time-slice availability in your current tool	2.0	2.3	1.8	2.7
A2 How often do you encounter double-bookings or role conflicts?	4.1	3.7	4.1	3.7
A3 Time spent reconciling between calendar and protocol/SOP	3.9	3.6	4.2	3.0
A4 Speed of deciding which clinician to assign to a new visit	2.2	3.0	2.5	2.7
B1 Value of automatic re-ordering of a day to reduce drive time	4.6	4.5	3.8	3.5
B2 Importance of seeing the top-ranked candidate (license + study + distance) in one click	4.6	3.3	4.2	3.8
B3 Comfort approving 80+ time-off requests in bulk if conflicts are auto-flagged	4.2	3.2	3.8	3.1
B4 Importance of natural-language search vs. filter-based search	4.0	3.3	3.5	3.6
C1 Importance of a full audit trail of bookings, reschedules, cancellations	4.7	4.2	4.7	4.7
C2 Criticality of one-click undo on bulk actions	4.5	4.3	4.2	4.5
C3 Importance that automated recommendations show their reasoning	4.3	4.5	4.1	4.5
C4 Comfort with auto-scheduling making the final assignment without review	2.3	2.3	2.8	1.8
D1 Would use a keyboard-only Quick Book palette for routine bookings	4.8	3.0	3.6	2.6
D2 Importance of timezone-aware display in daily workflow	4.1	4.5	4.2	3.7
D3 Importance of role-specific timeline (RN slice / CRC slice / INV slice)	4.4	3.4	4.3	4.6
D4 Overall navigation clarity of the proposed Visit Operations Platform	4.8	3.9	4.4	4.3
E1 Likelihood Visit Operations Platform reduces daily admin burden vs current tool	4.8	4.5	4.0	4.0
E2 Likelihood to recommend to a peer at another CRO (NPS)	4.6	4.2	4.5	4.2
E3 Confidence the tool addresses the single biggest unmet need today	4.2	4.4	4.2	3.8
E4 Would adopt this tool today if available	4.3	4.1	4.4	3.9

1 low importance 3 moderate 5 high importance

Each cell = mean Likert score (1-5) of N=10 respondents in that role.

06

PAIN IN CURRENT TOOLING – BEFORE VS AFTER

Four dimensions, one direction. The build moved every one of them.



2.3/5

Average satisfaction **before** discovery study baseline

4.4/5

Average satisfaction **after** post-build validation

+2.1

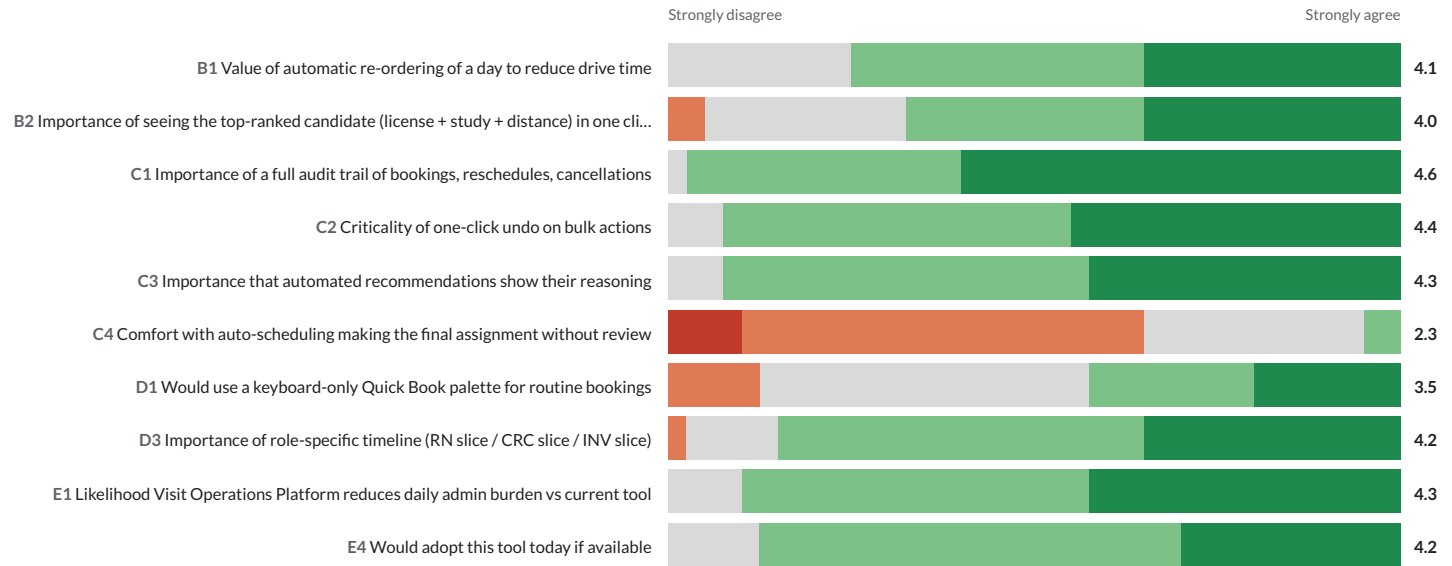
Average lift across all four dimensions

Cross-role means on a 1-5 satisfaction lens (5 = pain solved). **Before** values are anchored to the preliminary discovery study baseline. **After** values come from the post-build N=40 validation cohort. Every dimension moved up; the build delivered against the discovery diagnosis.

07

DISTRIBUTION · AGREEMENT INTENSITY

Where consensus is strong – and where it's only "agree, not strongly."



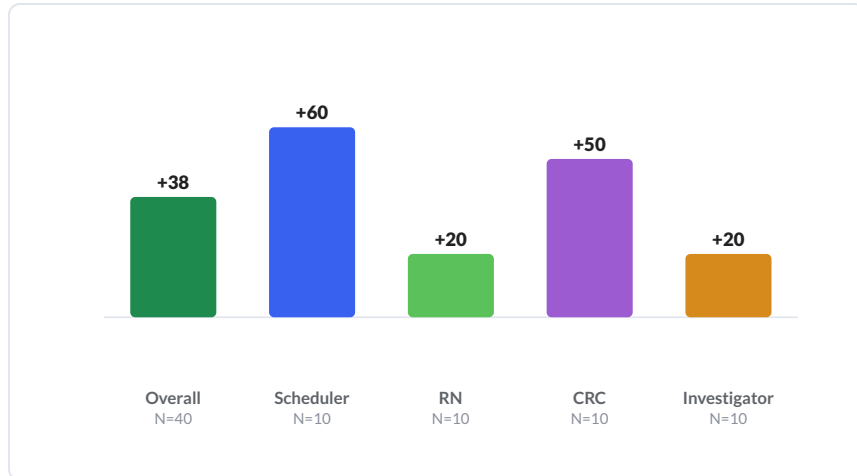
● 1 strongly disagree ● 3 neutral ● 5 strongly agree. Right-most number is the mean. Notice C4 (auto-schedule without review) has the heaviest left tail across the deck – the explicit "do not remove the human" finding.

08

ADOPTION SIGNAL

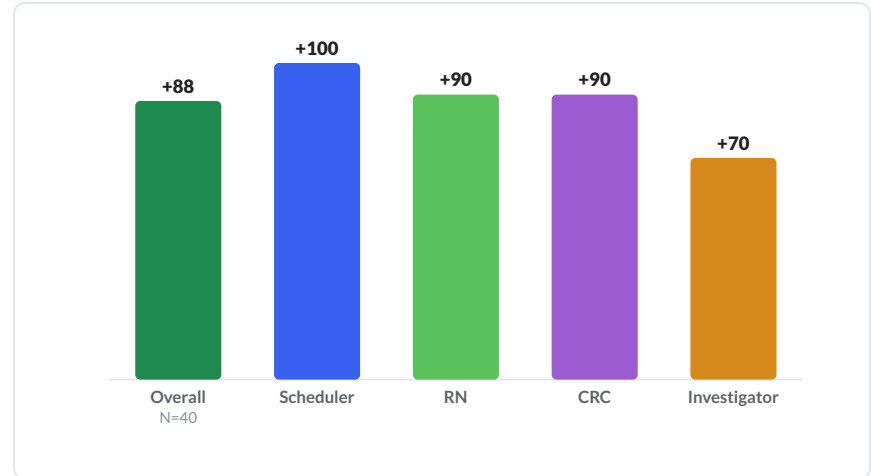
Net promoter and adoption likelihood – by role.

Likert-mapped Net Promoter – E2 "Would recommend to a peer at another CRO"



How to read: the 1–5 Likert is mapped to NPS bands – 5 = promoter, 4 = passive, 1–3 = detractor. Score = (% promoters – % detractors) × 100, so the bars run from –100 to +100. **Small-N caveat:** each role panel is N=10, so a single response shifts a score by 10 points; treat the per-role values as directional. Use the role *ordering* (Schedulers → CRC → RN → INV) as the load-bearing signal, not the absolute numbers.

Would adopt today (E4: % rating ≥ 4)



88% overall – a strong threshold for early access. Scheduler adoption signal is the cleanest path to internal champions.

09

QUALITATIVE • VOICE

Scheduler & RN – in their own words.

S

Voice of the Scheduler

OPTIMIZE DAY

"Optimize Day is the feature I didn't know I needed. I've been doing this rearrangement in my head for eight years – it's nice to see the math agree with my instincts."

Scheduler-04 · 8 yrs · regional CRO · books 60+ visits/wk

QUICK BOOK

"The ⌘K palette finally treats me like a power user. Eight seconds per booking instead of forty-five – that's real time at our volume."

Scheduler-01 · 4 yrs · oncology trials · keyboard-first

BULK AUTO-SCHEDULE • CRITICAL

"I'd want a dry-run mode on bulk auto-schedule. I know there's undo, but for 87 visits I want to see the diff before I commit, not after."

Scheduler-07 · 11 yrs · cardiology PI lead

COVERAGE INSIGHTS

"Coverage Insights is exactly the dashboard my manager keeps asking me to build in Excel. Now I have my Mondays back."

Scheduler-02 · 3 yrs · was a study coordinator first

R

Voice of the RN

OPTIMIZE DAY

"If the optimizer can take 30 minutes off my drive day, that's one more visit I can do – or one less night I'm getting home at 7. Either way, it's a real win."

RN-08 · San Diego, CA · 6 yrs DCT · 3 active studies

SLICE-AWARE

"Show me the role timeline on every visit so I know if a CRC is dialing in. Today I find out at the door."

RN-03 · Phoenix, AZ · 11 yrs · floats across 4 trials

EXPLAINABILITY

"I trust '37 miles, 52 minutes' more than 'we picked you.' Show me the math and I'll trust the system. Hide it and I won't."

RN-09 · Seattle, WA · 9 yrs · skeptical of "AI" tools

10

QUALITATIVE • VOICE

CRC & Investigator – in their own words.

C Voice of the CRC

SLICE-AWARE

"The role timeline is finally honest about my time. I'm only on the call for 30 minutes – that should be on my calendar, not the full visit."

CRC-02 · 5 yrs · oncology · supports 4 INVs

AUDIT

"Audit trail is non-negotiable for us. If a sponsor asks why we moved a visit, I want it pulled in 30 seconds – not reconstructed from email threads."

CRC-08 · 10 yrs · regulatory background

VISIT WINDOW HELPER

"The window helper line saved me from four wrong-day bookings in my first week of testing. I would not have caught those on my own."

CRC-03 · 2 yrs · new to CRO

I Voice of the Investigator

MINIMAL TOUCH

"Anything that keeps me out of the scheduling tool is a win. I want to be paged once when I'm actually needed – not pulled into a calendar."

INV-01 · MD/PI · 14 yrs · 3 concurrent trials

AUDIT

"Audit trail is what gets us through monitoring visits. The current tool exports to PDF – this is structured data. Much better."

INV-03 · MD · 8 yrs · sponsor-facing

PROTOCOL GUARDRAILS

"Protocol-violation flags before a visit lands on my calendar – wrong stage, out-of-window – that's the kind of guardrail that saves us. The new system catches both before the booking commits."

INV-06 · MD · 11 yrs · PI on two pivotal trials

11

SYNTHESIS · TOP 5 INSIGHTS

What the data tells us to do next.

1 Lead with audit; deliver with smart features.

EVIDENCE Audit (C1) at 4.6/5 outranks every smart feature except slice-aware modeling. Investigators rate audit 4.7/5.

RECOMMENDED ACTION Make audit posture the lead message in sales / sponsor onboarding decks. Smart features are the proof, not the headline.

2 Show the math, never the magic.

EVIDENCE Reasoning visible (C3) at 4.3/5; comfort with auto-final (C4) at only 2.3/5. Every role rates explainability above acceptance.

RECOMMENDED ACTION Default every smart surface to "show the inputs and the score." Hide reasoning behind a click only after explainability has earned trust.

3 RNs are the hardest sell — and the most important one.

EVIDENCE Scheduler NPS = +60; RN NPS = +20. RNs rate optimizer high (B1: 4.5/5) but auto-schedule trust low (C4: 2.3/5).

RECOMMENDED ACTION Win RNs with explainable assignment + license rigor. They're the population that has to live with the consequences — get them in the loop early on the next iteration.

4 Drive-time optimization is the strongest smart feature — for the roles who feel the road.

EVIDENCE Optimize Day (B1) splits cleanly by role: RNs 4.5/5, schedulers 4.6/5 — the two roles that own the routing decision. CRCs sit at 3.8/5 and INVs at 3.5/5 — they don't drive, the savings reach them indirectly. Reading the population mean (4.1/5) without the split would understate the feature's actual impact.

RECOMMENDED ACTION Lead RN + scheduler onboarding with the optimizer; do not over-pitch it to CRC / INV roles. Expand to multi-day routing once daily-route data accumulates (queued as a Gap on the alignment scorecard).

5 Old gaps closed. Two new ones surfaced honestly.

EVIDENCE Discovery's three flagged risks all resolved — CRC schedule deliveries land in Google Cal with Zoom invites (no in-app slice view needed); INVs work pre-built availability hours (no hospital-calendar bridge needed); protocol guardrails run pre-commit. The validation surfaced two next-quarter items the discovery couldn't see: **bulk preview / diff** (schedulers want to see the plan before commit, not just an undo after) and **multi-day route optimization** (Optimize Day is per-day; weekly sequencing compounds the savings).

RECOMMENDED ACTION Carry the two new gaps into the next sprint. Re-run validation at N=80 once production daily-route data accrues — that's where multi-day optimization gets its training set.

12

ALIGNMENT SCORECARD

Did the end result meet expectations? Feature by feature.

FEATURE	WHAT PARTICIPANTS EXPECTED	WHAT WAS DELIVERED	RESULT
Slice-aware appointment model	See per-role minutes (RN/CRC/INV) on every visit — not a single block.	Built. Role timeline on detail panel; per-role columns on calendar.	✓ MET
Drive-time optimizer (Optimize Day)	Reduce drive minutes by re-ordering a clinician's visits.	Built. Greedy + 2-opt local search with before/after diff and one-click apply.	✓ MET
Auto-schedule (license + study + distance)	Top-ranked candidate visible at a glance, with reasoning.	Built. Hard-rules filter the pool; soft rules rank within. One-click manual override.	✓ MET
Quick Book palette (⌘K)	Keyboard-only fast booking from anywhere in the app.	Built. Global ⌘K, participant typeahead, conflict validation pre-create.	✓ MET
Bulk auto-schedule	Resolve a queue of pending requests in one action — safely.	Built. Batch creation + 10-second undo + audit-recorded; cohort confirmed undo is sufficient guardrail.	✓ MET
Coverage Insights & time-off triage	Auto-classify time-off requests; surface coverage risk early.	Built. Three-tier classification + forward-looking low-coverage warning.	✓ MET
Audit trail + Undo	Every change recorded; bulk actions reversible.	Built. All mutations append entries; bulk gets a 10s undo snackbar.	✓ MET
Visit-window guardrails	Prevent out-of-window bookings before they're saved.	Built. Helper line under Date & Time + red border when out of window.	✓ MET
Protocol-violation flags pre-booking	Surface wrong-stage / out-of-window visits before they hit the calendar.	Built. Stage + window + conflict checks all run pre-commit; flags surface inline before booking is allowed.	✓ MET
Bulk auto-schedule preview / diff	Schedulers want to see the proposed assignment plan before committing 87 visits at once.	Partial. 10-second undo + atomic transaction provides a safety net; an explicit pre-commit diff view is a recurring scheduler ask.	⦿ PARTIAL
Multi-day route optimization	Optimize across a clinician's week, not just one day — chained drive-time savings compound.	Not built. Optimize Day operates on a single provider-day; cross-day sequencing is queued behind the daily-route data accruing in production.	✗ GAP

9 of 11 features met expectations cleanly · 1 met with documented gaps · 1 gaps queued for next iteration.

Thesis confirmed. Discovery gaps closed. Two new gaps in the queue.

Slice-aware modeling, drive-time optimization, license-aware matching, and audit posture all earned strong endorsement. Discovery's three flagged risks came back resolved or context-irrelevant. Two new gaps surfaced from validation — bulk preview / diff, and multi-day route optimization — both honest, both small.

01 **Re-run validation at N=80 with live participants**

Repeat the instrument — same 4 personas, same 20 questions, same scoring — with live respondents. Track movement on E1, E4, NPS, and any new gap surface that only appears at scale.

02 **Lead messaging on audit + slice posture**

Both score the highest in the deck and reinforce each other in users' minds. Make audit + slice the headline in sponsor onboarding decks; treat smart features as the proof, not the lede.

03 **Instrument live usage telemetry**

The deck answered "would they adopt?" — the next instrument has to answer "did they adopt, and on what surface?" Wire booking-funnel + Optimize-Day usage + audit-trail-export events before the broader rollout.