

RESEARCH REPORT

The staffing forecast report: how health system leaders predict gaps, use AI, and retain staff

What workforce leaders told us about predictive staffing, where AI has earned trust, and why flexibility is still the hardest lever to pull.

18

HOSPITAL & HEALTH
SYSTEM LEADERS

3

DISCUSSION THEMES

5

NEW INSIGHTS

EXECUTIVE SUMMARY

ABOUT THIS DISCUSSION

This report draws on a closed-door virtual roundtable held in August 2026 with workforce, nursing operations, and talent leaders from hospitals and health systems across the country. Eighteen leaders attended, spanning chief nursing officers, VPs of patient care, system staffing and scheduling directors, and other senior nursing operations executives. The format was deliberately informal, intended to be a guided conversation with three discussion blocks and participants spoke candidly about current-state operations. All comments here are de-identified by agreement.

The conversation was organized around three questions:

1. How do we move from reacting to anticipating?
2. Where can AI genuinely help, and where do we still need human judgment?
3. How do we design a workforce model people want to stay in?

Five patterns emerged across every conversation

01

Currently the Monday call-out is still solved by a person and a group text.

Across every model described, whether a unit manager, central staffing office, or house supervisor with a float pool, the last step is the same: broadcasting an open shift to an undifferentiated pool, with no sequencing by cost, familiarity, or fit. That gap gets paid for in premium rates and leader-worked shifts.

02

Three weeks is the forecast horizon leaders trust today.

Below 24 hours, one absence is manageable. Multi-day gaps need weeks of runway. Where that window exists today, one person rebuilds it by hand every week.

03

Float pools are being rebuilt around cost and proximity.

Centralized, premium-rate pools are giving way to campus-based flex at sustainable rates, with a small centrally held rapid-response group for day-of gaps. The unsolved question is which tier a given gap should draw from, and in what order.

04

So far AI is being used to inform, but not act

Forecasting, flagging, and screening are automated, with a human keeping the final say in every case. In every example discussed, AI functions as decision support by making recommendations, but it doesn't automate the underlying sequencing work.

05

Flexibility is the hard to deliver, but it doesn't have to be.

It means different things to different clinicians, which is exactly why it resists a single policy fix, and why it depends on knowing who your people actually are.

FINDING 1

The Monday call-out runs on people and text threads

Asked to walk through what happens when a clinician calls out on a Monday morning, participants described three recognizable models.

1

Unit-owned

The department manager or charge nurse takes the call and texts available staff directly.

2

Central staffing office

A dedicated staffer pages available staff and escalates to the manager if no one picks up.

3

House supervisor + float pool

The supervisor reallocates flex staff first, then the unit leader often ends up working the shift.

The models differ in who owns the decision, but they converge on the same final step: a human broadcasting to an undifferentiated pool, with no sequencing by cost, familiarity, or fit, and no way to know who will answer.

“That call-out gets turned into a text that goes out to all the available staff. And then we cross our fingers and hope somebody picks up a shift.”

NURSING LEADER, MULTI-HOSPITAL SYSTEM

A BETTER BROADCAST STILL ISN'T THE FIX

- Opt-in group messaging replaces phone calls, but is still a broadcast once a gap opens.
- Availability collected 1–2 weeks out gives a known list to work from.
- A leader covering the shift is a measurable signal of short-staffing.

THE ACTUAL FIX

Forecast the gap before it opens, then work a staffing waterfall automatically, internal float first, then flex pool, then external labor, so the most cost-effective tier is offered the shift first and fewer gaps ever reach the text-thread stage. That sequencing is what Medely's Talent Fusion automates.




The cost of the gap

Every hour an opening stays unfilled moves it up the incentive ladder, and the last resort, a unit leader working the shift personally, is the most expensive form of coverage there is, precisely because it never shows up in a labor line.

FINDING 2

Three weeks is the horizon leaders trust

A forecast is worth what it preserves in options. The earlier a gap is visible, the more it costs to ignore and the less it costs to fix.



ONE SYSTEM'S ROLLING REPORT: EVERY SUNDAY NIGHT, BY UNIT

One system uses color-coded red, yellow, and green by unit, and the bottom of the report shows how many flex pool clinicians are available for that campus in the same period. Leaders use it three ways: to release pre-scheduled incentives early instead of paying same-day emergency rates, to reallocate people across units before the week starts, and to recognize when the whole campus is red, meaning shared resources will not be available and the plan has to change.

“Having that glance at what the next three weeks looks like has been a game changer. We always look forward to getting it on Sunday.”

NURSING EXECUTIVE, MULTI-CAMPUS SYSTEM

Six-week scheduling cycles came up repeatedly as the other planning anchor. When a unit shows as consistently short across the six-week build, campus flex staff are pre-assigned to it rather than held back for day-of triage.

The practical lesson: a forecast is worth what it preserves in options. Two weeks of notice buys schedule redesign and lower-cost coverage. Two hours buys a premium rate. The three-week window is the point at which cost is still a choice.



The cost of the gap

A three-week forecast only works if it's maintained continuously and today, that usually means one person rebuilding it manually every week. Talent Fusion Optimize reads demand and staffing constraints continuously and generates that same three-week window automatically.

FINDING 3

Float pool design is being rebuilt around cost and proximity

BEFORE — PANDEMIC-ERA MODEL

A centralized, system-wide flex pool created quickly during the pandemic, paid at premium rates to stand it up fast. Effective at the time, but the rate structure made it unsustainable as a permanent model.



AFTER — EMERGING MODEL

Campus-based flex teams at sustainable rates, with staff working primarily at one campus and floating to other sites for critical needs. Sitting alongside that, a small group of flex clinicians is held centrally at the patient capacity center and deployed to any campus for same-day emergencies. The campus teams handle planned coverage; the central group absorbs the surprises. Leaders acknowledged losing some staff during the rate transition, and described the current model as steadier.

ALSO NOTED

A second system runs entity-level float pools alongside a system pool, with cost charged back to the requesting division. When float staff are not deployed, which is rare, in their experience, the time is redirected to education and unit support, so the investment stays productive. Making that cost visible at the divisional level changes how it is planned for

WHEN THE FORECAST SAYS SHORT, SOME SITES CHANGE THE MODEL INSTEAD

One ambulatory leader described shifting specific clinics to a telemedicine model on days when staffing is predicted to be tight, removing the need for on-site support staff rather than competing for them. Prediction creates options beyond filling the shift.



The cost of the gap

A tiered pool only pays off if each gap is routed to the cheapest tier that can actually cover it. Without that routing, organizations pay central rapid-response rates for gaps a campus team could have absorbed a week earlier, the same premium the pandemic-era model was dismantled to avoid.

FINDING 4

AI has earned trust in some areas. Automation is the next step.

One observation landed with visible agreement: AI is reading scans and drafting documentation at the bedside, while the decisions about who staffs that unit still run on spreadsheets and text threads. The clinical side of the house is adopting AI faster than the operational side.

Where AI is already being used among participants:



Demand & workload forecasting

Feeding an FTE forecasting algorithm, which in turn drives schedule prediction. Built first for inpatient, emergency, and obstetrics; ambulatory is next.



Admin load reduction

Flagging that a clinician has not met their FTE commitment or that a shift is missing a charge nurse, checks no leader has time to perform manually across a 300-person department.



Day-of decision support

A seven-day look-ahead that incorporates incentive policy and historical call-out patterns, including the fact that Mondays run heavier.



Workforce planning

Modeling seasonality and historical turnover so a unit with a known October or February vacancy wave can build bench strength ahead of it.



Clinical inbox support

Drafting responses to patient messages to offload volume from nurses and physicians.



Candidate screening

One organization piloted an AI voice agent for first-round nursing assistant screens. Availability at any hour cut the drop-off caused by business-hours-only scheduling, while talent acquisition and the unit leader kept the decision.

"AI has created visibility faster on areas of opportunity to improve throughput. We use AI as an accountability tool to create and support our processes.

-NURSING OPERATIONS EXECUTIVE

FINDING 4 CONTINUED

Where the market stands today on AI



Every leader using AI for scheduling described a suggestion engine, not an auto-assigner, a starting point, not a ceiling. The opportunity is to automate everything that doesn't require judgment; matching, sequencing, executing the staffing waterfall and only bring a leader in when a decision actually calls for one.

That is the gap Optimize is built to close: automating forecasting, matching, and waterfall execution end-to-end, and surfacing a decision to a leader only when one is genuinely required, rather than gating every shift on manual review.



Adoption is change management

As one participant put it, innovation that people do not trust simply does not get used. One caution repeated by multiple leaders: AI output is only as good as the underlying data. When information is scattered across spreadsheets, text threads, and multiple vendor systems, the recommendations cannot be relied on.

How trust gets built

Asked whether accuracy or explainability matters more for trust, one leader named a third thing: a visible commitment that humans still run daily operations. Data-driven insight is there to give leaders time back, with everyone knowing a person who cares about the team and the patients is still making the call.

FINDING 5

Flexibility is the currency and the hardest one to deliver

Asked what keeps clinicians who have stayed for years, participants listed convenience and location, working the shift they want, their coworkers and the patients they serve, culture, and growth that is visible rather than promised. Flexibility is hard to deliver because it means different things to different people.

Schedule input

A say in their own hours and shift choices.

Specialty movement

Moving between units by choice; one leader described a clinician cross-skilled across several units who moves between them by choice to avoid burning out in any one.

Predictability

Consistent staffing levels were themselves named as a retention lever, because repeatedly stretching the same people accelerates the exits that make staffing worse.



CONTRACT LABOR MIX AFFECTS CORE STAFF SATISFACTION

A unit can be fully staffed on paper and still feel unstable to permanent staff when continuity is low, in cases where teams can't reliably bring back the same trusted clinicians, even if the shifts are technically covered.


VIRTUAL NURSING EXTENDS CAREERS

Moving admissions, transfusion checks, and similar tasks to virtual roles lowers the load on the floor, and it keeps experienced nurses contributing and mentoring when 12-hour shifts at the bedside are no longer sustainable for them.



The cost of the gap

Continuity and known clinicians are the same missing capability, knowing who your people are. Talent Fusion holds that knowledge in the system rather than in individual memory, surfacing internal staff, float, and repeat clinicians a unit has already worked with, including favorites, so continuity becomes something you can schedule for instead of hope for.



\$60,090

Avg cost to replace an RN

FINDING 6

Fragmentation is quietly costing you



GEOGRAPHIC

One organization's two closest facilities sit nearly 100 miles apart. Floating staff isn't realistic, prediction matters more, not less, when the response options are narrower.



DATA

Vendor sprawl was named as a self-inflicted problem, more systems make data harder to pull. Consolidating partners is a data strategy as much as a purchasing decision.



The cost of the gap

Fragmentation sets the ceiling on every other finding here. A forecast is only as good as the record behind it, and a waterfall can only sequence the tiers it can see. Each additional vendor system is one more blind spot the forecast has to work around.

WHAT LEADERS CAN ACT ON NOW

1

Audit your staffing path.

Count how often a unit leader ends up working the shift, the clearest available proxy for structural short-staffing.

What it requires: a coverage record that captures who actually filled every shift, not just that it was filled.

2

Replace the manual outlook with an automated forecast.

A rolling three-week, color-coded view with flex capacity, generated automatically instead of rebuilt by hand.

What it requires: a forecasting layer like Medely's Optimize that generates and refreshes the three-week window automatically, so every leader has it without the manual lift.

3

Separate planned coverage from day-of response.

Campus-based flex teams for predictable gaps, a small central group for emergencies, clear economics for each.

What it requires: costed visibility into every tier and automated routing between them, so the cheapest adequate tier is offered the shift first.

4

Set the AI guardrail before the pilot.

Recommendations to leaders, final say with humans, and a parallel run to earn credibility.

What it requires: a system that automates sequencing and execution end-to-end while assignment authority stays with a person.

5

Define what flexibility means for each population.

Schedule input, multi-specialty movement, and virtual roles solve for different clinicians, not one initiative.

What it requires: clinician-level knowledge — specialty, history with the unit, stated availability, held in a system rather than in individual memory.



SEE HOW TALENT FUSION CLOSES THIS GAP

Every finding in this report converges on the same two gaps: leaders cannot see far enough ahead, and when they do see a gap, there is no automated way to work it down the cost curve.

Talent Fusion closes both, with a single forecasted view across internal staff, float pools, and flexible coverage, and a waterfall that offers each shift to the most cost-effective tier in order. That is the difference between planning on a three-week window and reacting on a same-day one.

SAME-DAY



3-WEEK WINDOW

Reactive. Premium rate. One option.

Forecasted. Waterfall-ready. Full options.

[Schedule a demo of Talent Fusion](#)