APSIS SOLUTIONS — EHR OPTIMIZATION RESOURCES

The EHR Optimization Playbook

Workflow Redesign · Adoption Recovery · Reporting Stabilization · Post-Live Governance

Why go-live is the starting gun, not the finish line — and what the first 18 months after cutover actually require.

There is a phrase that circulates in EHR implementation circles, usually said with a tired laugh: “Pick the right software and you can forget about it.” Every organization that has lived through the eighteen months after a go-live knows that sentence is exactly backwards. Cutover is not the finish line. It is the starting gun for the work that actually determines whether an EHR becomes a tool clinicians trust or a system they’ve learned to route around.

The period right after go-live — informaticists sometimes call it the “messy middle” — is where physicians are still mentally living in two systems: the one they knew and the one they’re now required to use. They feel the new EHR was something done to them rather than something built with them. That feeling doesn’t resolve itself. It either gets addressed through deliberate, workflow-specific optimization, or it calcifies into the workarounds, the after-hours “pajama time,” and the slow specialty-by-specialty erosion of trust that turns an EHR into the thing everyone complains about at every meeting for the next five years.

This playbook is built around four disciplines that determine the difference: workflow redesign that reflects how clinicians actually practice rather than how the build was originally specified, adoption recovery for the specific behavioral patterns that predict burnout and attrition, reporting and dashboard stabilization so the data layer doesn’t quietly degrade while everyone is focused on clinical workflow, and the governance structure that prevents the next eighteen months from looking like the last eighteen.

1.2 hrs
average daily ‘pajama time’ on clinic days
Journal of Internal Medicine, 2025

#1
cause of physician burnout
Tebra 2025 Physician Burnout Survey

88,758
alert firings/month for a single retired feature
Centra Health, 2024 Epic alert audit

103 hrs
annualized savings from optimizing 200 order sets
Carilion Health physician-builder program

CHAPTER 1 THE MESSY MIDDLE

Why the post-go-live period is the most consequential — and the most misunderstood — phase of any EHR program

Cutover Is an Event. Adoption Is a Process.

Most EHR project plans treat go-live as the deliverable. The build is validated, the data is migrated, the training is complete, the command center stands down after hypercare, and the project team moves on to the next initiative. From a project management standpoint, the box is checked.

From a clinical standpoint, nothing has actually been proven yet. The build was validated against the workflows the design team specified eighteen months earlier — not against how a busy cardiology clinic actually runs on a Tuesday afternoon when three patients are double-booked and the attending is covering for a colleague who called in sick. The training was generic enough to get every user through go-live, but it was not specific enough to teach an orthopedic surgeon how their chart review, their preference lists, and their in-basket should be configured differently than a primary care physician’s. The data layer was tested for go-live readiness, not for what happens when six months of real clinical volume starts generating reports nobody validated against production scale.

Every one of those gaps is invisible at go-live. They become visible over the following twelve to eighteen months — as rising help desk ticket volume, as physicians who quietly build personal workarounds instead of raising issues, as dashboards that drift out of alignment with what departments actually need, and as the EHR satisfaction surveys that show declining scores in the exact specialties whose workflows were least represented in the original build.

The trust deficit no implementation plan accounts for

Physicians experience an EHR implementation as something that happened to them. The build decisions were made in conference rooms by IT staff and a handful of physician informaticists representing dozens of specialties they don't personally practice. By the time a neurologist or an orthopedic surgeon sits down at the new system, decisions that directly shape their daily workflow were made without their input — and they know it. Optimization is the process of giving that input back, after the fact, in a way that actually changes the system. Skip this step and the trust deficit becomes permanent: clinicians stop reporting problems because they've concluded nothing changes when they do.

The Three Patterns That Predict Optimization Failure

- **organizations that fold post-live issues into the general IT ticket queue — competing for prioritization against printer problems and password resets — are signaling that EHR experience is not a priority. Optimization needs its own governance, its own prioritization criteria, and its own dedicated resourcing separate from break-fix support.** Treating optimization as a backlog instead of a program:
- **a workflow fix that helps primary care can actively hurt a surgical specialty using the same order set with different defaults. Optimization that isn't segmented by specialty workflow — chart review habits, in-basket structure, documentation patterns, order set composition — produces "average" improvements that don't move the needle for the specialties generating the most tickets.** Specialty-blind optimization:
- **the single most corrosive pattern in post-live EHR programs is the survey that goes out, generates feedback, and produces no visible change. Clinicians who give feedback that disappears into a black hole stop giving feedback — not because they're satisfied, but because they've learned it doesn't matter.** No mechanism for clinicians to see that their feedback changed something:

CHAPTER 2 WORKFLOW REDESIGN

How specialty-specific workflow analysis turns generic builds into tools clinicians actually want to use

Why 'How the System Works' Is the Wrong Training Question

The shift that separates effective post-live training from go-live training is a shift in the question being answered. Go-live training answers "how does the system work?" — necessarily generic, because everyone needs to function on day one regardless of specialty. Post-live optimization answers a completely different question: "how do you work, and how does the system need to be configured to support that?"

The documentation needs of an orthopedic surgeon are not the documentation needs of a primary care physician — not just in what data gets captured, but in how they search the chart, how they manage task messages, how they interact with the in-basket, and how follow-up care gets structured. A workflow redesign program that doesn't segment by these differences produces generic improvements that help everyone a little and nobody enough to change their experience of the system.

The Workflow Audit: What Gets Reviewed, By Whom, In What Order

STEP	FOCUS AREA	WHAT GETS REVIEWED	SUCCESS SIGNAL
1	**Order Set & Preference List Review** *Owner: Physician-builder team + pharmacy*	Audit the highest-volume order sets and medication preference lists by specialty. Identify defaults that don't match how the specialty actually prescribes — wrong default doses, missing commonly-paired orders, order sets built for inpatient use that surgical clinics are using off-label. Carilion Health's physician-builder team reworked preference lists across 200 commonly ordered medications and produced an annualized savings of 103 hours of physician time.	**Reduced order entry clicks per encounter; fewer ad-hoc order modifications**
2	**Alert & BPA Governance Review** *Owner: Clinical informatics + safety committee*	Pull a month of alert firing data and review against actual override and action rates. Alerts with extremely high firing rates and low action rates are training clinicians to click through warnings reflexively — which undermines the alerts that actually matter. Centra Health found one retired feature generating 88,758 alert firings per month; silencing it eliminated nearly 90,000 unnecessary clicks without removing any clinical safeguard.	**Alert override rate trending down for retained alerts; total alert volume reduced**
3	**In-Basket & Inbox Structure Review** *Owner: Nursing informatics + specialty champions*	Review in-basket categories, pooled vs. individual routing, and message volume by type for each specialty. The "Inboxologist" model — a structured triage role for routing and resolving in-basket messages before they reach the physician — has shown meaningful reductions in primary care message burden in recent published work.	**Reduced physician-facing in-basket volume; faster message resolution time**
4	**SmartPhrase & Documentation Template Review** *Owner: Specialty champions + clinical documentation team*	Review documentation templates and SmartPhrase usage by specialty. Templates built for the average use case create extra clicks for specialties with materially different documentation patterns — a dermatology visit note and a behavioral health intake note should not be using the same structural template with minor field changes.	**Reduced note-writing time per encounter; increased SmartPhrase adoption in target specialties**
5	**Chart Review & Navigation Pattern Review** *Owner: Clinical informatics + Epic Signal analysis*	Use Epic Signal data on chart review time, navigator usage, and tab/flowsheet configuration to identify specialties spending disproportionate time hunting for information. Nested order panels have shown modest but measurable ordering efficiency improvements for high-utilizer primary care physicians in recent published research.	**Reduced chart review time per encounter; fewer navigation clicks to common destinations**

The Physician-Builder Model

The organizations getting the most out of post-live optimization are not the ones with the largest IT teams. They are the ones who have built a physician-builder program — clinicians with Epic build certifications who remain active in clinical practice. The advantage isn't just technical access. It's that a physician-builder is

both the person identifying the workflow problem and the person who understands, from direct daily experience, what "fixed" actually looks like.

A non-clinical analyst can build exactly what's requested and still produce something that doesn't work in practice, because the request didn't capture a nuance that only becomes obvious when you're the one using it during a patient encounter. A physician-builder closes that gap. Carilion's team described it directly: "What's great about the physician builder team is, we can improve clinical workflow efficiency, because we're not just builders, we're also users of the system."

What this means for organizations without an internal physician-builder program

Not every organization — particularly mid-size health systems and ambulatory groups — has the volume to justify dedicated physician-builder FTEs. The alternative that works: structured workflow shadowing sessions where a certified Epic analyst spends time directly observing specialty-specific clinical workflows (not interviewing about them — observing them), paired with a rapid build-and-validate cycle where changes are tested with the same clinicians who identified the problem before being rolled out broadly. The principle is the same even if the staffing model differs: the person making the build decision needs direct exposure to how the workflow actually happens.

CHAPTER 3 ADOPTION RECOVERY

Reading the signals that predict burnout, attrition, and quiet abandonment — before they become resignations

What Epic Signal Can and Cannot Tell You

Epic Signal generates a rich set of efficiency metrics: time in the EHR, pajama time, time outside scheduled hours, chart review time, in-basket time, and dozens of others. Organizations frequently make two opposite mistakes with this data. The first mistake is ignoring it entirely and relying on anecdote. The second mistake — increasingly common as Signal dashboards become more prominent — is treating Signal data as a diagnostic tool that can identify which specific clinicians are burned out or at risk of leaving.

KLAS's Arch Collaborative research, conducted in partnership with Epic, has been direct about this: Signal data is not a meaningful or predictive measure of whether a specific provider is dissatisfied, burned out, or considering leaving. Two physicians can show identical after-hours EHR time, and one is perfectly happy with their workflow while the other is one bad week from resignation. The data shows correlation at a population level, not prediction at an individual level.

What Signal data is genuinely useful for: identifying **where** to focus optimization effort — which specialties, which workflows, which clinics show efficiency patterns meaningfully different from peers — and measuring whether an intervention changed those patterns afterward. It's a population-level prioritization tool and a before/after measurement tool. It is not a clinician risk score, and treating it as one erodes trust if clinicians learn their EHR usage is being individually monitored as a proxy for their wellbeing or commitment.

The Adoption Recovery Metrics That Matter

METRIC	HEALTHY RANGE	WARNING PATTERN	SOURCE	WHAT TO DO ABOUT IT
Pajama time (after-hours EHR work)	**Trending flat or down month over month, by specialty**	Sustained >1 hour/clinic day with no improvement after training interventions	*Epic Signal pajama time metric*	Workflow review for the specific specialty — documentation templates, order sets, in-basket structure

Alert override rate	**High-severity alerts: low override rate. Low-severity: lower volume overall**	Override rate climbing across the board — alert fatigue, clinicians no longer reading any alert	*Epic alert/BPA reporting*	Alert governance review: retire or reposition low-value, high-volume alerts
Help desk ticket volume by department	**Stable or declining post-hypercare, proportional across departments**	One or two departments generating disproportionate ticket share weeks after go-live	*Help desk / ServiceNow reporting*	Targeted workflow audit for the specific department — this is the highest-signal indicator of a build gap
In-basket message volume per provider	**Within normal range for specialty; pooled messages being triaged before reaching physician**	Rising individual in-basket volume with no corresponding pooling/triage structure	*Epic Signal in-basket metrics*	In-basket structure review; consider triage/'inboxologist' role for highest-volume specialties
EHR satisfaction survey trend (Arch Collaborative or equivalent)	**Stable or improving year over year, consistent across specialties**	Declining trend concentrated in specific specialties not represented in original build team	*KLAS Arch Collaborative / internal survey*	Specialty-specific workflow review prioritized by lowest satisfaction trend, not just lowest absolute score

The Feedback Loop That Rebuilds Trust

If physicians have concluded that giving feedback doesn't change anything, the single highest-leverage intervention available to an optimization program is making a visible, attributable change based on feedback — and communicating that the change happened because of what was reported.

This sounds simple and is routinely skipped. The mechanism: maintain a visible "optimization log" — even a simple one — that specialty champions and department leads can see, showing recently submitted issues, their status, and what changed when they were resolved. The Centra Health alert reduction, the Carilion preference list work — these become trust-building events specifically because they were communicated back to the clinicians whose feedback drove them, not just implemented silently in a build update nobody heard about.

CHAPTER 4 REPORTING & DASHBOARD STABILIZATION

The data layer doesn't announce when it breaks — it just quietly stops being trusted

Why Reporting Drift Is the Silent Optimization Failure

While workflow redesign and adoption recovery are happening in plain view — surveys, tickets, visible training sessions — the reporting and analytics layer is drifting in the background, and almost nobody notices until someone asks why two dashboards disagree.

The pattern is consistent: at go-live, the critical reports were built and validated against the test environment. Six months later, clinical volume has grown, new departments have come online, build changes from the optimization work in Chapters 2 and 3 have altered the underlying data in ways the original reports didn't anticipate, and nobody re-validated the reporting layer against any of it. The report still runs. It still returns a number. The number may no longer mean what everyone assumes it means.

The connection between workflow optimization and reporting integrity
Every workflow change made in Chapters 2 and 3 has a downstream reporting consequence. If a department changes how it documents discharge disposition to fix a workflow pain point, every report that aggregates discharge disposition — readmission dashboards, length-of-stay reports, quality measure calculations — needs to be checked against the new documentation pattern. Optimization programs that treat workflow redesign and reporting validation as separate workstreams consistently produce 'optimized' workflows that broke a dashboard nobody was watching.

The Post-Optimization Reporting Checklist

- **any change to documentation templates, order sets, or encounter type assignments that fed Chapter 2's workflow redesign needs a corresponding check against the Clarity and Caboodle reports that consume that data.** Re-run report validation after every workflow build change:
- **reports that ran acceptably at go-live volume can degrade significantly as patient volume, documentation volume, and report subscriber count all increase over the following year. A report that took 90 seconds at go-live and now takes 12 minutes is a signal that the underlying SQL and indexing strategy needs review — not that the organization simply "outgrew" the report.** Confirm Clarity ETL performance hasn't degraded as clinical volume has grown:
- **the eighteen months after go-live is when dashboard sprawl begins. Departments that requested a dashboard during build, used it for a few weeks, and then stopped — because the workflow it was meant to support changed — leave behind a dashboard that nobody is maintaining or validating, but that still appears in search results when someone is looking for "the readmission report."** Audit dashboard usage and retire what isn't being used:
- **when workflow optimization changes how a specialty documents something, and a new report gets built to track the impact of that change, confirm its metric definition matches the certified definition used elsewhere in the organization. This is precisely how organizations end up with two readmission rates that don't agree — one built before optimization, one built after, using slightly different logic.** Reconcile metric definitions across the reports that came out of optimization work:

CHAPTER 5 OPTIMIZATION GOVERNANCE

The structure that prevents the next eighteen months from looking like the last eighteen

Optimization Is a Program, Not a Project

The organizations that sustain EHR satisfaction gains treat optimization as a permanent operating function with its own governance, not a one-time project that wraps up with a final report. The governance structure doesn't need to be large — even a lean structure with clear ownership outperforms a large team without clear accountability.

The Minimum Viable Optimization Governance Structure

ROLE	RESPONSIBILITY	CADENCE
Optimization Owner (CMIO or delegate)	Final prioritization authority for build changes. Balances competing specialty requests. Owns the optimization log and	Weekly review of open items; monthly report to executive sponsor

	communicates closed-loop feedback to clinicians.	
Specialty Champions	One per major service line. Surfaces workflow issues specific to their specialty, reviews proposed build changes before they go live, communicates changes back to their colleagues.	Bi-weekly check-in; ad hoc for urgent issues
Build/Analyst Team	Implements approved changes. Validates downstream reporting impact (Chapter 4) for every workflow change before it goes live.	Sprint-based; 2–4 week build cycles
Data/Analytics Lead	Owns the post-optimization reporting checklist. Confirms metric definitions remain consistent across reports affected by build changes.	Reviews every build release for reporting impact
Executive Sponsor	Removes organizational barriers, resolves cross-specialty prioritization conflicts that champions can't resolve at their level, ensures optimization isn't deprioritized when other initiatives compete for the same build resources.	Monthly governance review

Where AI Documentation Tools Fit — and Where They Don't Replace This Work

Ambient listening and AI-assisted documentation tools are a genuinely significant development in EHR optimization, and recent published data backs that up: one health system study found average note-writing time dropped from 6.58 minutes to 5.45 minutes within three months of an ambient listening pilot — with note length increasing, meaning the time savings weren't coming from less thorough documentation.

But these tools are an addition to optimization work, not a substitute for it. An AI scribe that transcribes a clinical encounter into a note still has to insert that note into a chart whose navigation, in-basket structure, and order workflows are exactly what they were before. If the underlying workflow problems described in Chapters 2 through 4 aren't addressed, an AI documentation layer makes the documentation step faster while leaving every other friction point in the EHR experience untouched. Organizations that treat an AI scribe rollout as "the optimization initiative" and stop other workflow work are solving one part of the problem and calling it done.

The measure of optimization success isn't a satisfaction score. It's a trend.
A health system that goes live with average EHR satisfaction and sees that score hold steady or improve over the following two years has an optimization program working as intended. A health system that goes live with strong satisfaction and watches it erode specialty by specialty over the same period has a build that was never revisited after the implementation team moved on. The difference between those two outcomes is almost never the underlying software. It's whether anyone kept doing the work — the workflow audits, the alert governance, the reporting validation, the closed-loop feedback — after the go-live banners came down.

and physician re-engagement — deployed in days, not quarters.

About Apsis Solutions Inc.

Apsis Solutions Inc. is a Tampa-based healthcare IT consulting firm specializing in Epic implementation and optimization, Clarity/Caboodle analytics, IT workforce staffing, cybersecurity compliance, and managed database services. We run structured post-live optimization programs — workflow redesign, adoption recovery, reporting stabilization, and governance design — for health systems whose go-live is behind them but whose optimization work never started, stalled, or quietly drifted. We serve health systems, academic medical centers, federal agencies, and prime contractors as a capable prime vendor, strategic subcontractor, and managed services partner. The agility of a boutique firm. The capability of a prime. — apsisinc.com | 813-999-4853 | info@apsisinc.com