Credo

Closing the Data-to-Action Gap in Value-Based Care

Value-based care (VBC) succeeds when the right clinical actions happen reliably at the point of care, and when the documentation supporting those actions **accurately reflects patient complexity and performance under contract.**

Physicians are the lynchpin in that system. They make the clinical decisions that drive utilization and outcomes, and they produce the documentation that substantiates acuity and performance. Any technology—including AI—should therefore be evaluated by how well it optimizes physician engagement: how quickly it delivers relevant context, how much it earns trust through supporting evidence, how seamlessly it fits into day-to-day workflow, and how clearly it informs and supports the physician's recommended course of action.

AI doesn't replace clinical judgment. Its role is to **amplify physician expertise** by reducing the search-and-synthesis burden, improving confidence in the longitudinal record, and making decisions and documentation easier to execute under real time constraints.

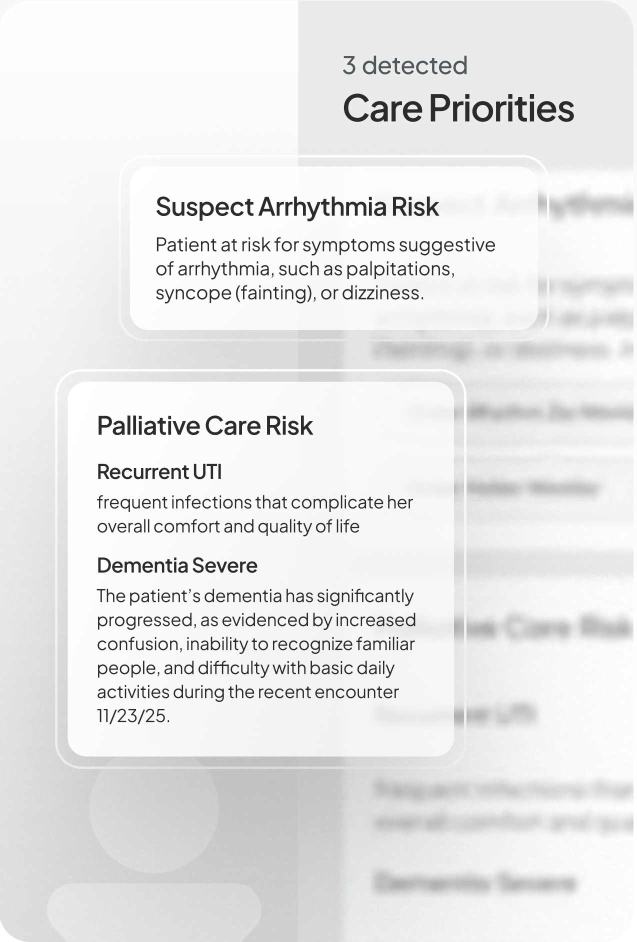

3 detected

Care Priorities

Suspect Arrhythmia Risk

Patient at risk for symptoms suggestive of arrhythmia, such as palpitations, syncope (fainting), or dizziness.

Palliative Care Risk

Recurrent UTI

frequent infections that complicate her overall comfort and quality of life

Dementia Severe

The patient's dementia has significantly progressed, as evidenced by increased confusion, inability to recognize familiar people, and difficulty with basic daily activities during the recent encounter 11/23/25.

Table of Contents

Executive Summary

VBC doesn't fail because organizations lack insight. It fails when insights don't arrive fast enough, or in a form clinicians can use to drive the required action at the point of care.

Two realities define whether a VBC program succeeds:

01 **Clinical action** (diagnoses, orders, referrals, medication changes, follow-up plans) drives utilization, outcomes, and total cost of care.

02 **Clinical documentation** (accurate, specific) determines whether risk-bearing entities can validate patient acuity and performance under contract.

This creates a hard truth: **physician engagement is the delivery mechanism for value.** Even the best data will underperform if clinicians can't access complete context, verify evidence quickly, and translate insights into defensible decisions and documentation inside a time-boxed encounter. This paper synthesizes:

→ Clinician survey results on point-of-care friction.

→ Learnings from a 150+ clinician beta programs focused on designing point-of-care tools that connect longitudinal data to physician workflow.

→ Credo operational findings on record completeness and downstream impact.

→ Foundational research on missing clinical information in ambulatory care.

The Baseline Problem: Missing Clinical Information in Ambulatory Care

The most frequently cited data from JAMA on the issue of missing patient records quantified how often clinicians walk into visits without the full picture of a patient's health history. Across **1,614** primary care visits, clinicians reported missing clinically important information in **13.6%** of encounters.[1] The most common missing elements were lab results, letters/dictation from other clinicians, radiology results, history and physical exam data, and medication information. [1]

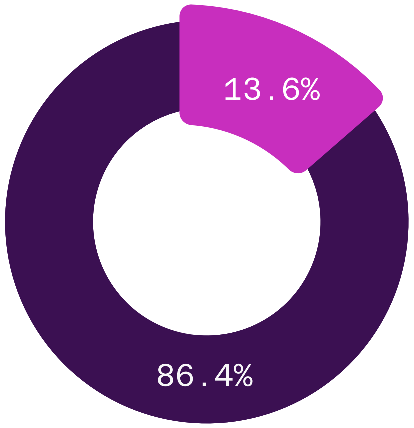

More Than 1 in 8 Clinical Encounters Lack Critical Information

Encounters with missing information

Common missing elements:

- Lab results
- External clinician notes/dictation
- Radiology results
- H&P data
- Medication information

Encounters without missing information

Clinicians also reported:

- **52.3%** of missing information was located outside their clinic's clinical system (but within the U.S.).[1]
- Missing info was at least somewhat likely to adversely affect patient well-being in **44%** of missing-info visits.[1]
- Missing info was likely to cause delays or additional services in **59.5%** of missing-info visits.[1]

Bottom Line

Physicians believe at least one-in-eight ambulatory care visits occur without the records clinicians believe they need, creating delays, duplicate work, and avoidable risk.[12]

The Data: Missing Records Missing in Ambulatory Care

The JAMA data on this issue quantified just how often clinicians walk into visits without the full picture. Across 1,614 primary care visits, clinicians reported missing clinically important information in 13.6% of encounters. The most common missing elements were lab results (6.1% of all visits), letters/dictation (5.4%), radiology results (3.8%), history and physical exam data (3.7%), and medication lists (3.2%). (NCBI)

What Credo's Data Shows

93.5% of targeted patients have additional records outside owned datasets

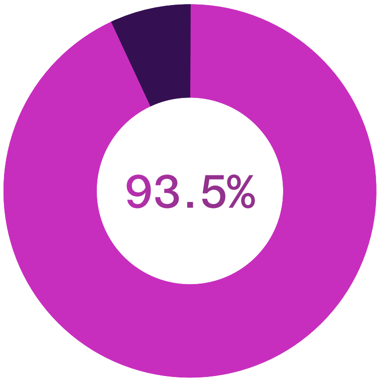

What Clinicians Think

13.6% of patients missing data

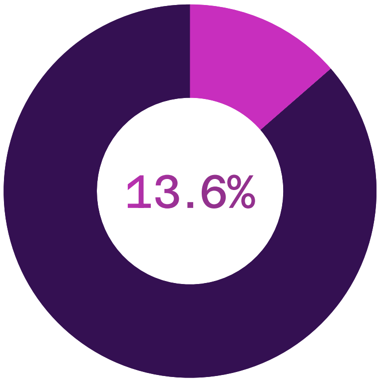

The Data

Continued

Metric	Value
Total primary care visits observed	**1,614**
Visits with clinically important information missing	**13.6%**
Missing info “outside our clinical system but within the US”	**52.3%**
Missing info thought at least “somewhat likely” to adversely affect patient well-being	**44.0%**
Missing info expected to cause delay in care or additional service(s)	**59.5%**

Impact Dimension	% of missing info visits (n=220)
Missing info was outside the clinic’s clinical system (but within the US)	**52.3%**
Missing info at least somewhat likely to adversely affect patient well-being	**44.0%**
Missing info likely to cause delay in care or additional services	**59.5%**

Missing Records Type	% of missing-info visits	% of all visits (n=1,614)
Lab results	45.0%	61%
Letters/dictation	39.5%	5.4%
Radiology results	28.2%	3.8%
History & physical exam	26.8%	3.7%
Current & prior medications	23.2%	3.2%
Pathology results	15.0%	2.0%
Immunization records	12.3%	1.7%
Procedures	7.3%	1.0%
Other	5.0%	0.7%

Clinician time spent looking unsuccessfully	% of attempted visits	% of all missing-info visits
5–10 minutes	**25.6%**	14.5%
> 10 minutes	**10.4%**	5.9%

The Cost of Missing Records: Operational, Clinical, and Financial Impact

Missing records are not merely a data quality issue. They constrain clinical capacity, increase patient safety risk, and in value-based care, undermine contract performance. The impact is most visible at the point of care, where clinicians must make decisions in real time.

01 Modeled Operational Impact

Credo Health modeled the impact of missing/late records as an attribution problem across three domains: clinical operations, clinical harm, and risk-based revenue. **Assumptions are documented in the appended** and grounded in published prevalence of missing records during visits and Credo Health's observed record incompleteness.

Missing or late records drive missed visits, late cancellations, reschedules, and "redo" appointments, especially in referrals and specialty care where prerequisites (outside notes, imaging, labs, discharge summaries) determine whether the visit can proceed.

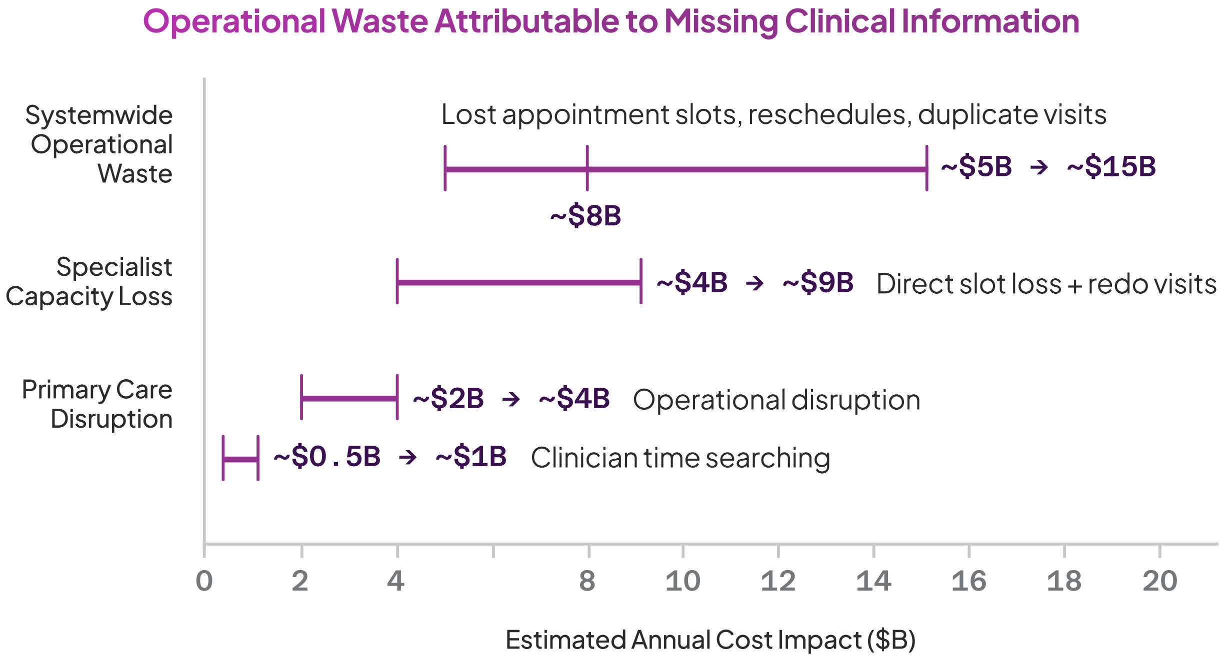

Operational Waste Attributable to Missing Clinical Information

02 Modeled Clinical Impact

When clinicians lack timely access to prior context, key results, and referral documentation, care is delayed, repeated, or delivered with incomplete information. In aggregate, missing record access is estimated to contribute to:

~10,000 – ~30,000
Deaths per year

~1.0M → ~2.5M
Delayed or missed diagnoses

Note: These are attribution estimates, not claims that the literature directly reports "X deaths caused by missing records." They combine national diagnostic error burden work with evidence that missing information is common and disruptive at the point of care.[12]

03 Modeled Financial Impact in Risk-based Contracts

In Medicare-risk populations, incomplete longitudinal records lead directly to under-captured acuity and missed performance. The modeled opportunity from making records "complete enough" and converting evidence into reliable suspecting and documentation is:

- **~$1,100 per patient per year** in risk-based revenue/ premium opportunity (range ~$800–$1,900), depending on contract structure and realization.[7]

Bottom Line

Missing records break the chain from **data to trust to action**. In real-world use cases, Credo Health sees AI effectively addressing breaks in this chain, when it strengthens physician engagement by making longitudinal context more complete and usable, as well as easier to verify and act on.

TEFCA: The Regulatory Framework Designed to Close the Gap

The Trusted Exchange Framework and Common Agreement™ (TEFCA) sets a common framework for sharing information across multiple "exchange purposes," including Treatment, Individual Access, Payment, Health Care Operations, Public Health, and Government Benefits Determination. Qualified Health Information Networks (QHINs) were established with the expectation that they would support these purposes, and today the Exchange Purposes SOP requires responses for Treatment and Individual Access Services. (ASTP TEFCA RCE)

Health Information Exchanges (HIEs) across the United States

Under the Trusted Exchange Framework and Common Agreement (TEFCA), state and regional Health Information Exchanges (HIEs) act as critical "on-ramps" that connect local healthcare providers to the broader nationwide network.

Currently, there are approximately 75 to 100 operational HIEs across the United States, a number that has consolidated over the last decade. While the terms are often used interchangeably, a state HIE typically covers a single state's entire jurisdiction—often with state government backing or designation—whereas a regional HIE serves a specific geographic market that may cover just a portion of a state or span across multiple state lines.

Under TEFCA, these organizations are intended to act primarily as intermediaries (known as Participants or Sub-participants) rather than national backbone networks (QHINs). They are used to aggregate data from local practices, hospitals, and public health agencies and connect them to a Qualified Health Information Network (QHIN). This structure ensures that complex local needs—such as specific patient consent laws, social service data, and precise patient matching—are managed at the community level while still enabling nationwide interoperability.

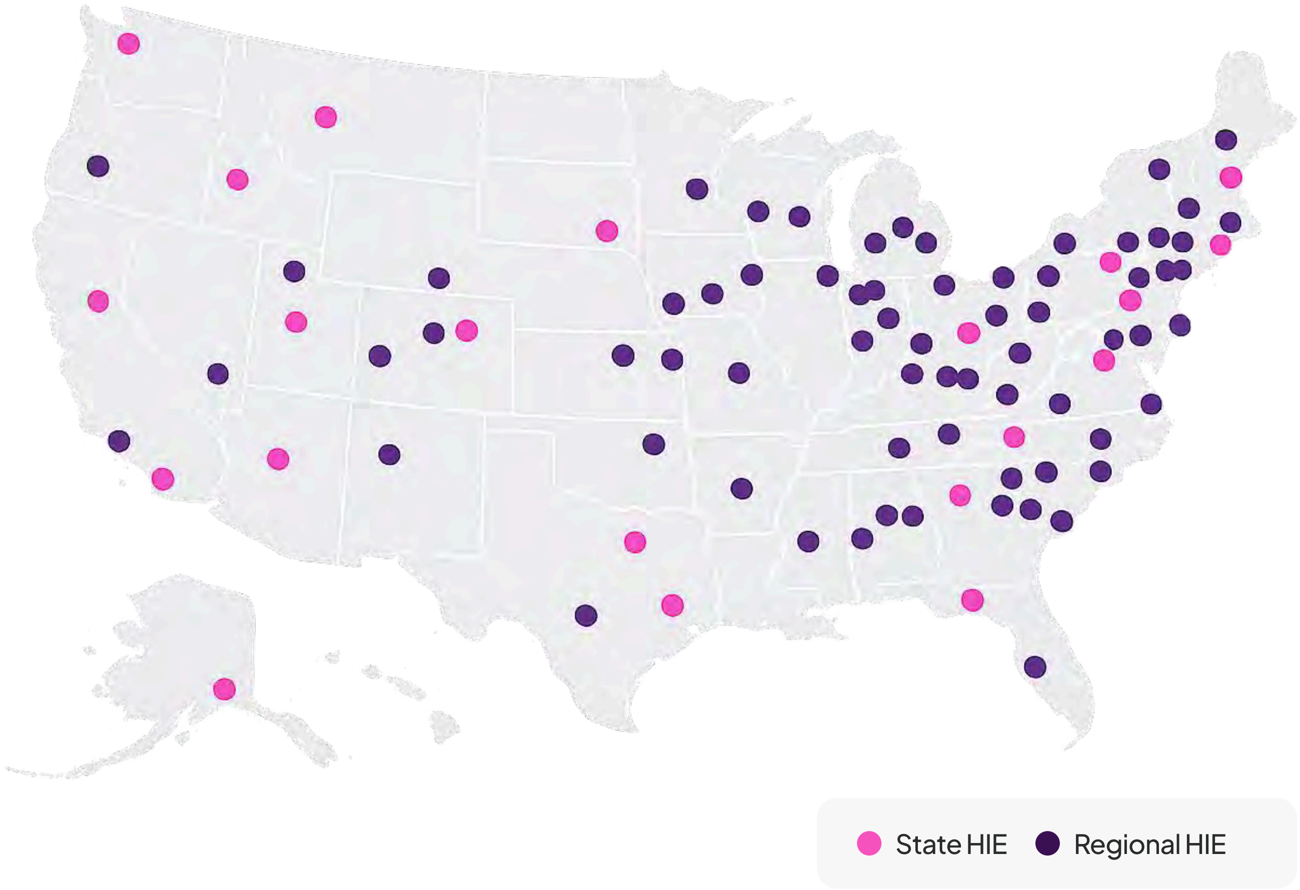

Why Interoperability Isn't Enough at the Point of Care

While TEFCA was designed to make health data exchange more consistent and trustworthy nationwide, the framework alone has not automatically produced better record completeness at the point of care. In practice, three operational realities still break the promise of "one connection equals the whole chart":

- **Coverage isn't uniform:** Not every site of care participates equally, and not every clinically relevant artifact is available through a single feed.
- **Clinical history is multi-modal:** Key evidence lives across narrative notes, PDFs, imaging reports, and consult letters; often in inconsistent data structures.
- **Workflow still matters:** Even when data can be retrieved, it may not arrive in time or in a form physicians can trust and use to meet the physician's need for a time-boxed visit.

This aligns with national interoperability measurement: ONC reports that in 2023, 70% of non-federal acute care hospitals engaged in all four domains of interoperable exchange (send, find, receive, integrate) routinely or sometimes, while fewer routinely do so.[10]

The practical point-of-care reality is that "exchange exists" does not mean "the longitudinal record is always complete enough and usable in workflow."

What Credo Sees in Practice:
The Longitudinal Record Is Still Incomplete

Interoperability has improved the plumbing, but the "complete record" remains more aspiration than reality. Even when an ACO, MSO, health system, or payer believes it has a robust dataset (claims, an HIE feed, EHR data, etc.), a meaningful portion of the patient story frequently lives elsewhere.

> **Credo's Care Map findings:** most patients have additional records outside owned datasets

HCCs (Hierarchical Condition Categories) are diagnosis groups CMS uses to measure patient complexity and set risk-adjusted payment in Medicare Advantage—so when supporting documentation lives outside the "owned dataset," organizations can systematically under-measure acuity. In 25 customer pilots, Credo used Acquire Digital AI to retrieve records across QHINs and connected HIEs, while Care Map identified gaps and triggered targeted chases.

Across all pilots, Credo consistently found that “owned” datasets missed meaningful parts of the record:

- Across pilots, Credo retrieves an average of 63.4 records per patient for 93.5% of patients (results vary)[11]
- Across broader deployments, Credo retrieves an average of 52 records per patient for 90%+ of patients (results vary)
- Across Credo customers and pilots, filling record gaps consistently uncovers ~1 to 1.5 net-new HCC per patient (typically higher for HCCs count for new patients vs established patients.[11]

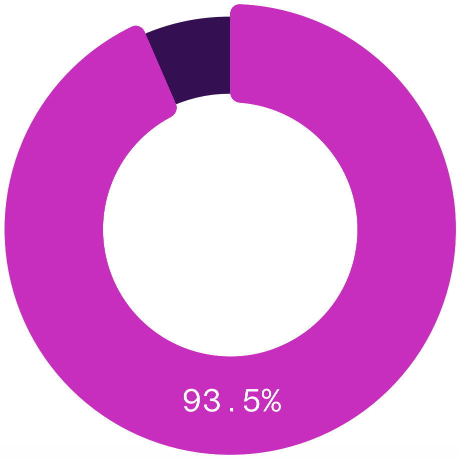

More Than 1 in 8 Clinical Encounters Lack Critical Information

93.5%

of patients have additional records outside owned datasets

63.4

Clinical documents per patient in targeted searches

~1 to 1.5

Net-new HCC per patient

What this means: for many value-based populations, a meaningful share of the patient story is not accessible through a single feed or “owned dataset.” Closing that gap requires a hybrid approach: digital retrieval plus targeted, AI-assisted chase for what doesn’t arrive electronically.[11]

Why the full data set is as important as the AI “insights”

This completeness gap is the hidden root cause behind downstream problems:

- Physicians don’t trust summaries if they suspect the record is incomplete.
- Risk and quality programs struggle to operationalize next actions when key evidence is missing.
- Even the best point-of-care experience fails if the underlying record lacks critical consults, diagnostics, and transitions of care.

In Summary

Completeness is the prerequisite for consistent, trustworthy synthesis and physician action.

The Point of Care Bottleneck in 2025:

What Has (and Hasn't) Changed

In Fall/Winter 2025, Credo Health worked with and surveyed 400+ physicians and staff across conferences, onboarding calls, and online surveys. The dominant barriers were still **time**, **completeness**, and **access**:

- **46.5%** cited time constraints during the visit as a top barrier to using more of the available data.
- **45.6%** pointed to incomplete patient records.
- **35.1%** reported difficulty obtaining a complete patient history.
- **29.9%** specifically called out lack of access to previous records.
- **12.4%** highlighted the burden of reviewing visit history across encounters.[12]

Side-by-side comparison: external research vs. Credo surveys vs. Credo chart retrieval

Theme / Measure	External research	Credo surveys	Credo Chart Retrieval (Acquire + Care Map)
Core problem: clinicians reporting missing clinically important information	13.6% of visits missing important info[1]	45.6% view incomplete records as top barrier[12]	In pilots, additional records exist for 93.5% of patients[11]
Access gaps / "not all the history is here"	52.3% of missing info outside system[1]	35.1% difficulty obtaining complete history; 29.9% lack access[12]	52 records are identified per patient on average[11]
Time friction during encounter	Extra searching time common when info missing[1]	46.5% cite time constraints as top barrier[12]	Hybrid retrieval closes gaps pre-visit; targeted chase varies by source responsiveness[11]

The throughline: ambient tools reduce *documentation* friction, but VBC still struggles to reliably deliver complete, usable longitudinal context and make it fast and verifiable enough for clinicians to act.

> **Key implications for value-based care:** the physician engagement gap isn't just "missing information." It's the persistent inability to reliably turn scattered history into decision-ready, source-verifiable clinical context within a time-boxed encounter so that the desired action actually happens.

Proof the Chain Closes:
2.1× More Net-New Accepted Suspects at the Point of Care

Completeness and synthesis only matter if they translate into clinician trust and action at the point of care. In a head-to-head pilot with a large value-based primary care organization, Credo's physician-trained coding AI was evaluated against a competing industry solution using the same ~500-patient cohort provided by the organization. Both vendors produced a similar volume of suspects, but surfaced meaningfully different candidates, showing that the real test is not how many suspects you generate, but how many clinicians can verify and accept.

Results (same cohort, side-by-side):

- Suspects identified: ~1,000 each (only ~10% overlap)
- Net-new suspects: Credo ~550 vs competitor ~300
- Acceptance rate (net-new): Credo 77% vs competitor 66%

Net-new accepted by physicians:

- Credo: ~425 (+112%)
- Competitor: ~200

Why this matters: the difference wasn't just "more suspects." It was more accepted suspects, the outcome that actually drives documentation quality and risk capture, enabled by (1) more complete encounter history and (2) evidence-backed suspects clinicians could verify quickly and easily. In other words, Credo closed the last mile from data to trust to action in a measurable way, delivering 2.1× more net-new accepted suspects from the same patient population.

The “Documentation to Action” gap compounds in Medicare populations

Medicare beneficiaries with multiple chronic conditions are the heaviest users of care, and therefore the most exposed to fragmentation. CDC analysis highlights that two-thirds of beneficiaries with **2+ chronic conditions account for 93% of Medicare spending, and one-third with 4+ account for almost three-fourths of spending.** (CDC)

Meanwhile, the number of “handoffs” that can strand information keeps rising. CMS notes that from **2019 to 2000**, the share of Medicare beneficiaries seeing five or more physicians annually increased from 18% to 30%, and primary care providers must coordinate with far more clinicians than in prior decades. (CMS)

> **Implication for VBC:** the patients who matter most financially and clinically are also the patients most likely to have critical records scattered across sites of care, making point-of-care completeness and synthesis even more consequential.

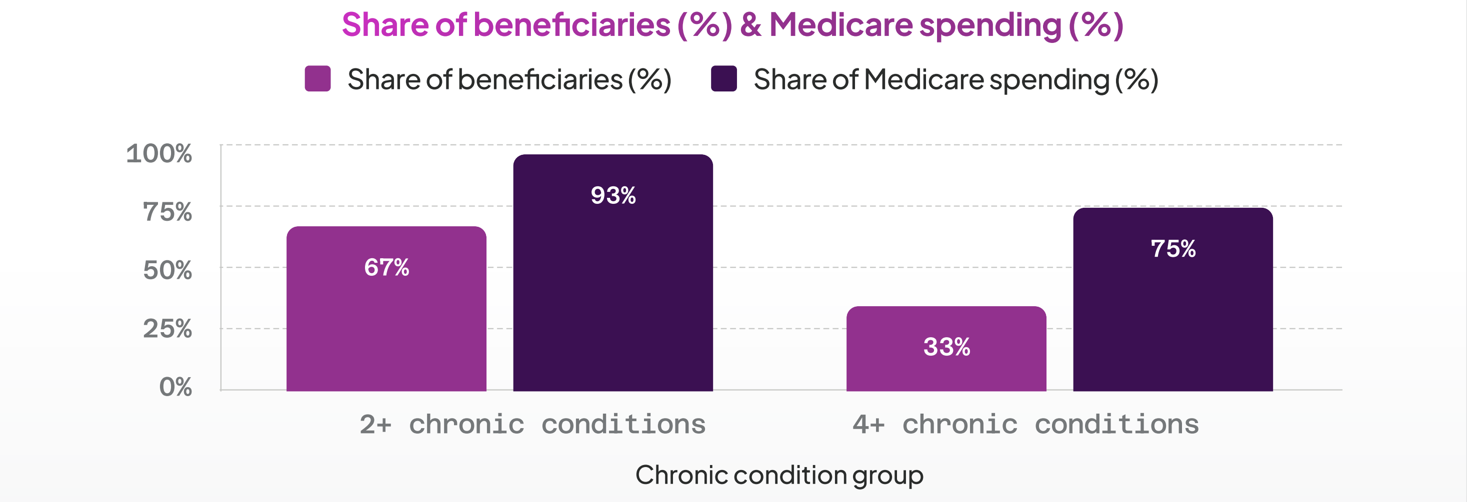

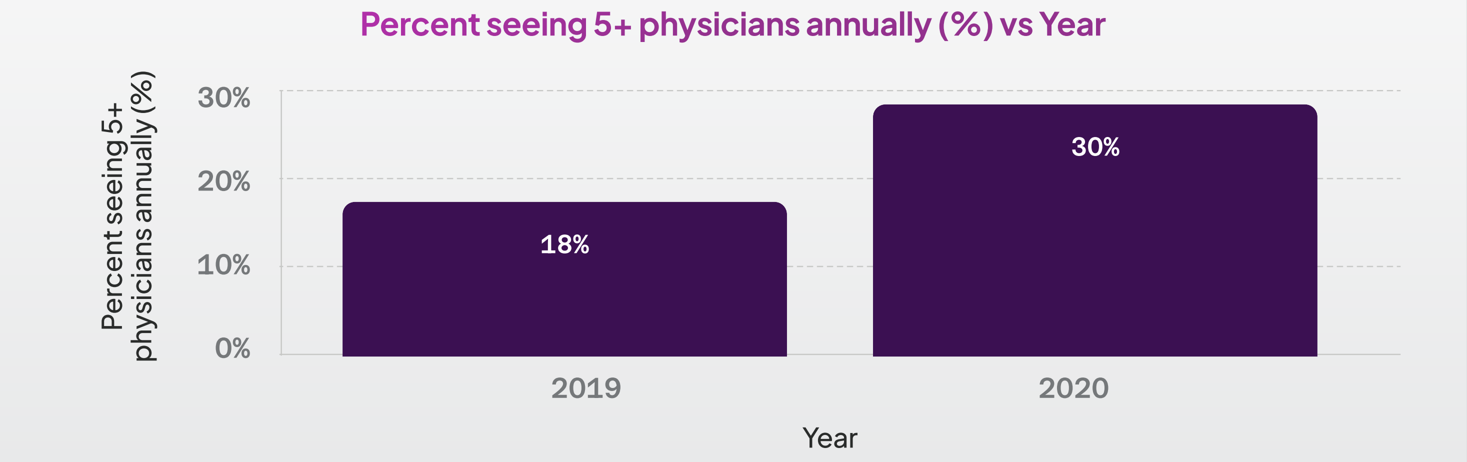

The Capacity Constraint: Clinicians Are Already at the Limit

Point-of-care engagement competes with a simple constraint: time.

EHR event-log research found family medicine physicians spent **355 minutes/day** (5.9 hours) in the EHR per weekday per clinical FTE, including **86 minutes/day** after clinic hours.[16] Meanwhile, the AMA reports that physician burnout remains high (e.g., **48.2%** reporting at least one symptom of burnout in 2023), reinforcing that added friction at the point of care reduces adoption headroom.[17]

The Time Constraint at the Point of Care

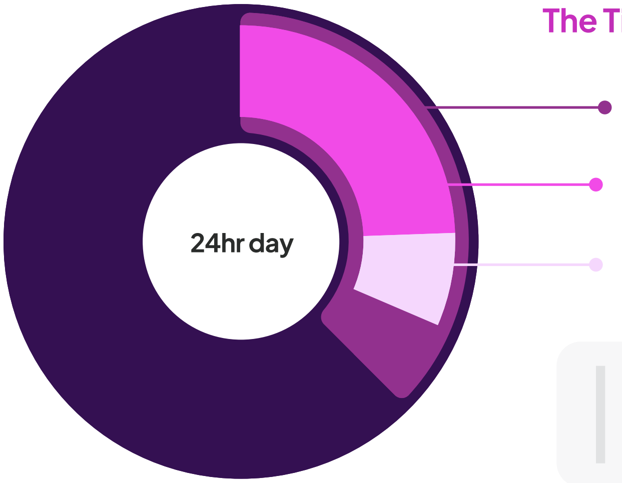

48.2%

of physicians report burnout symptoms

So the bar for any VBC "insight" isn't just "is it true?" It's whether it is:

- Fast enough to use during the visit
- Source-verifiable (clear provenance)
- Clinically relevant to the decision in front of the physician
- Immediately usable (actionable without extra hunting)

In Summary

Interoperability is able to move data, but VBC requires that data to land in the workflow in a way that enables action.

Compilation of Physician Burnout Data from National Surveys

Theme	Metric	Value
Burnout	Physicians reporting **≥1 symptom of burnout**[1]	**48.2%**
Burnout trend	Physicians reporting **≥1 symptom of burnout**[2]	**45.2%**
EHR time load	Total EHR time per weekday[3]	**355 min/day (5.9 hrs)**
“Pajama time”	EHR time **after clinic hours**[3]	**86 min/day (1.4 hrs)**
Workweek + composition	Average physician workweek and components[4]	**5hrs/wk** total; **27.3** direct care; **14.1** indirect care; **7.9** admin
After-hours EHR distribution	Physicians spending **>8 hrs/week** on EHR outside normal hours[4]	**20.9%**
	Physicians spending **0–2 hrs/week** on EHR outside normal hours[4]	**~26%**
	Physicians spending **6–8 hrs/week** on EHR outside normal hours[4]	**14%**

1) AMA: 2023 “Organizational Biopsy” national comparison report. 2) Stanford Medicine: 2024 National physician survey series. 3) Source: Annals of Family Medicine; Dr. Adam Stewart (142 family physicians). 4) AMA: 2023 aggregated report results

The Insights: Credo Insights from a 150+ Clinician Point-of-Care Beta

Credo enrolled 150+ clinicians in a point-of-care beta focused on one objective: helping excellent clinicians move faster from longitudinal data to confident action inside real visits.[18]

What we built and tested: a digital chart acquisition tool paired with a clinically trained AI insights engine designed to identify and extract clinically relevant data from health records, plus a clinically trained LLM that can answer physician questions about patient history with source-cited evidence (so trust can be verified, not assumed).[18]

Across the beta, friction patterns showed up in ~25% of physician questions/queries as moments where clinicians needed more context, clearer evidence, or a more usable output to act quickly.[18]

01 Physicians don't need to ask one question; they need to build query stacks.

Clinicians reduce uncertainty by layering evidence over time. A typical stack looks like:

- "What happened during the last hospitalization?"
- "What were the last three A1Cs with dates?"
- "Any nephrology notes?"
- "Current meds, plus start dates and prescribers?"

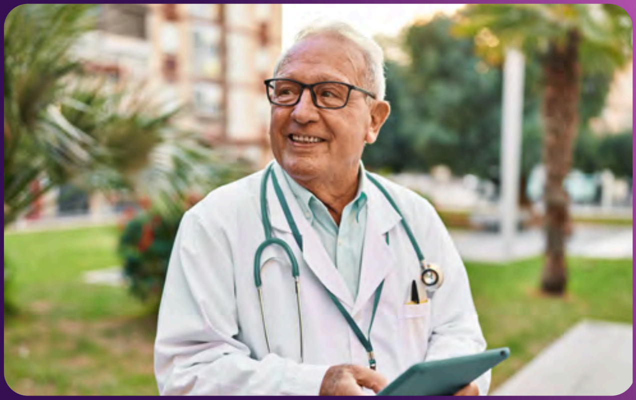

> **What Credo AI enabled:** working with physicians, Credo engineers compressed multiple stacked questions into a single hypothesis-level query (e.g., "Is this patient a good candidate for a GLP-1?"), returning a source-cited aanswer that compiles the building blocks physicians would gather manually.[18]

02 Bundling increases speed, but only if every building block is source-cited.

Clinicians repeatedly told us:

- "Show me where that came from."
- "Let me click into the evidence."
- "If I can't see the notes/labs, I can't use it."

> Bundling works when the system doesn't just summarize but also shows its work.[18]

03 “Clinically relevant” beats “comprehensive” and timelines often become the missing layer.

Clinicians didn’t ask for more pages. They asked for:

- What’s unstable
- What supports documentation without extra scavenger hunts
- What changed
- What’s missing that could change today’s decision

> In many friction moments, clinicians weren’t just looking for facts. Often they were looking for patterns (trends, escalation/de-escalation, recurrences). Timelines provided the organizing structure that allowed faster judgments under time pressure.[18]

04 Great answers aren’t the finish line. Turning answers into usable artifacts is the next value link.

The most-used outputs were the ones clinicians could quickly adapt into:

- HPI / interval history
- problem-oriented assessment inputs
- prior-auth support
- risk/quality documentation with clear evidence

> Insight is necessary, but compounding value comes when outputs become artifact-ready without forcing clinicians to reformat, re-hunt, or re-prove what’s already in the chart.[18]

A Practical Playbook to Increase Physician Engagement at the Point of Care

`Pillar 1`

Make the record "complete enough" before the visit

Records need to be complete enough to reduce surprise and rework:

- Ensure recent discharges, consult notes, key diagnostics, and medication changes are discoverable
- Close gaps that alter decisions (missing labs/imaging/med list provenance)

`Pillar 2`

Convert record volume into decision-ready summaries

Because clinicians can't read everything:

- Emphasize timelines, deltas, and why it matters
- Keep everything traceable to source evidence (click-through)

`Pillar 3`

Operationalize query stacks as reusable bundles

Turn common stacks into one-click panels:

- "New patient snapshot"
- "Diabetes + kidney risk panel"
- "Recent utilization and transitions"
- "Cardio diagnostics panel (echo/ECG/cath + key labs)"

`Pillar 4`

Measure engagement like a product metric, not a training outcome

Track:

- Time-to-answer
- Artifact reuse (copy/adapt into documentation)
- Source-open rate (trust behavior)
- Follow-on actions (orders, referrals, gap closure signals)

Then feed that back into iteration with clinicians continuously.

Specialty Deep Dive: What Frontline Clinicians Asked For at the Point of Care

We worked closely with frontline physicians and care teams using Chart w/GPT in real workflows. The goal was simple: understand what earns trust and saves time when the record is finally "complete enough," and what still blocks action inside a 15–20 minute visit.[18]

Who we heard from

Feedback came from outpatient and inpatient contexts, including:

- **Wound care / home-visit support teams**
- **Advanced Heart Failure Cardiology**
- **Neurology** (diagnostics-heavy)
- **Oncology**
- **Pediatrics**
- **Infectious Disease** (clinic + inpatient consult workflows across multiple inpatient EHRs)
- **Reproductive Endocrinology/IVF**
- **Integrative / complex chronic care**
- **Practice operators and administrators**

Across specialties, the pattern was consistent: clinicians don't want "more data." They want clinically relevant, source-grounded answers that fit how they already think and chart.[18]

Feedback themes that repeated everywhere:

> ***"Make the outside packet usable in seconds."***
> Clinicians don't read 80–120 page packets end-to-end, they hunt for the same artifacts every time: last note, key diagnostics, current meds, and what changed.[18]

> ***"Let me search and synthesize like a clinician, not an EHR."***
> Clinicians want tools that support stacking and bundling, not one-off answers.[18]

> ***"If you bundle the answer, show every building block."***
> Trust becomes brittle when provenance isn't obvious.[18]

> ***"A chronological diagnostics timeline I can chart from."***
> For many specialties, a clean chronological diagnostics list is the fastest path to action.[18]

> ***"Make it real inside my workflow, copy and push-to-note."***
> Clinicians consistently asked for a direct path from insight to documentation artifacts.[18]

Conclusion

Closing the Last Mile from Data to Trust to Action

VBC performance depends on a simple chain: **data to trust to action**.

Physicians are the only actors who can reliably complete that chain at the point of care. AI creates value when it **optimizes physician engagement** by making the longitudinal record more complete, transforming it into decision-ready context, grounding it in evidence clinicians can verify, and packaging it into artifacts that fit real workflows.

> ### Missing records are not just a data problem
>
> Missing records are a capacity problem, a patient safety problem, and a contract performance problem. Closing the last mile requires making the record "complete enough," converting complexity into decision-ready, source-verifiable context, and delivering outputs that support physician action under time pressure.

Appendix: Modeling Note Summary

Credo Health modeled the impact of missing/late records as an attribution problem across three domains—operations, clinical harm, and risk-based revenue—grounded in published prevalence of missing information at visits and Credo's observed record incompleteness. Operational waste is estimated by combining national ambulatory visit scale with evidence that ~10–15% of outpatient encounters occur with clinically important information missing, and that ~32% of missing-information visits experience a delay or disruption of care (used as an anchor for the disruption rate), then applying conservative assumptions for the fraction that translates into lost appointment capacity, duplicate/redo visits, and staff/clinician time spent searching for missing information (calibrated to published time-to-search distributions). Clinical harm estimates are modeled as a conservative attributable share of national diagnostic error burden, using evidence that missing clinical context is common and can delay or disrupt care; these figures are directional and not causal counts. Financial impact in risk-based contracts is modeled by converting incremental evidence found in patient records (as measured in Credo retrieval cohorts) into risk-based revenue/premium opportunity, with realization ranges reflecting variability in documentation quality, contract structure, and operational adoption.

Appendix A: Modeling Notes + Source Map for the Impact Estimates

A1. Operational waste and capacity loss

(~$8B/year total; specialist ~$4B–$9B; primary care ~$2B–$4B + clinician time).

**These are Credo modeled estimates designed to translate "missing information is common and disruptive" into operational magnitude. They are not directly published as "$X due to missing records."*

Source map used to ground the model inputs:

- Volume baseline: **~1.0B physician office visits/year**; **50.3%** to primary care (CDC FastStats).[19]
- Missing-info prevalence and disruption: **13.6%** missing important info in primary care; missing info often outside system and associated with delay/additional services (JAMA 2005).[1]
- Independent confirmation of missing info rates and disruption/harm signals in outpatient specialty context (NHS outpatient study: 15% missing; 32% delay/disruption; 20% risk of harm).[2]
- No-show environment calibration: systematic review average **~23%** no-show rate across studies (used for context, not as "records cause no-shows").[20]

Credo modeling approach: Estimate the fraction of visits where missing records cause (a) reschedule/cancellation, (b) redo visit, and (c) staff/clinician time waste; multiply by a conservative per-visit operational cost and capacity value. Credo keeps this conservative because the literature rarely isolates "records caused it" cleanly.[3]

A2. Deaths and delayed/missed diagnoses

(~20,000 deaths/year; ~1.5M diagnoses)

These are attribution estimates that triangulate national diagnostic error burden with evidence that missing info is common and disruptive at the point of care.

Source map:

- Outpatient diagnostic errors: **~5.08%** of U.S. adults; **~12 million** adults/year (BMJ Quality & Safety).[5]
- Serious harms from diagnostic error: national estimate including **~371,000 deaths/year** and **~424,000 permanent disabilities** (BMJ Quality & Safety / Johns Hopkins summary).[6]
- Missing information is common and linked to delays/disruption and perceived risk (JAMA 2005; Burnett 2011).[12]

Credo modeling approach (high level): apply a conservative attribution fraction ("missing/inaccessible records meaningfully contributed") to national error/harm totals; choose a range that reflects uncertainty and avoids overstating causality.[4]

A3. Risk-based contract opportunity

(~$1,100 PMPY; range ~$800–$1,900)

This combines:

- Credo retrieval evidence that net-new records exist for most patients (93.5%) and that meaningful HCC value is missed in standard coding (0.815/patient).[11]
- A modeled conversion from incremental captured risk to PMPY premium/revenue opportunity that varies by contract structure and realization.[7]

Footnotes

01 Smith PC, Araya-Guerra R, Bublitz C, et al. Missing Clinical Information During Primary Care Visits. JAMA. 2005;293(5):565–571. JAMA Network | PubMed (index)

02 Burnett SJ, Deelchand V, Franklin BD, Moorthy K, Vincent C. Missing clinical information in NHS hospital outpatient clinics: prevalence, causes and effects on patient care. BMC Health Services Research. 2011;11:114. HTML | PDF

03 Credo internal modeling memo (Dec 2025): operational waste + capacity loss attributable to missing/late records, calibrated using national visit volumes and missing-info prevalence/disruption studies, with conservative attribution assumptions. (Internal; not publicly cited.)

Footnotes

Continued

04 Credo internal modeling memo (Dec 2025): attribution of missing/inaccessible records as a meaningful contributor to fatal harm and delayed/missed diagnoses, triangulating diagnostic error burden with missing-info prevalence/disruption evidence. (Internal; not publicly cited.)

05 Singh H, Meyer AND, Thomas EJ. The frequency of diagnostic errors in outpatient care: estimations from three large observational studies. BMJ Quality & Safety. 2014;23(9):727–731. BMJ Q&S article | Europe PMC full text

06 Newman-Toker DE, Wang Z, Zhu Y, et al. Burden of serious harms from diagnostic error in the USA. BMJ Quality & Safety. 2024;33(2):109–120 (journal issue pagination may vary by format). BMJ Q&S article page | BMJ Q&S full PDF

07 Credo internal modeling memo (Dec 2025): PMPY risk-based opportunity using Credo net-new record prevalence and HCC value signals plus modeled conversion to revenue/premium, varying by contract.

08 ASTP / Sequoia Project RCE. Exchange Purposes Explained (TEFCA)

09 ASTP / Sequoia Project RCE. Exchange Purposes (XPs) Implementation SOPs (resource library). (General TEFCA & RCE resource library)

10 Office of the National Coordinator for Health IT (ONC). Interoperable Exchange of Patient Health Information Among U.S. Hospitals: 2023 (Data Brief).

11 Credo Health internal data (Acquire + Care Map), 2025: aggregated cohort across ~25 customers and thousands of MA patients (owned dataset vs net-new records; net-new artifacts; average HCC value identified). (Internal; not publicly cited.)

12 Credo Health internal surveys (Fall/Winter 2025): 400+ physicians and staff; barriers to point-of-care data use. (Internal; not publicly cited.)

13 Credo internal head-to-head pilot: ~500 patients; suspect overlap; acceptance rates; net-new accepted HCCs. (Internal; not publicly cited.)

14 Centers for Disease Control and Prevention (CDC). Prevalence of Multiple Chronic Conditions Among Medicare Beneficiaries, United States, 2010. Preventing Chronic Disease. 2013. HTML | PDF

15 Centers for Medicare & Medicaid Services (CMS). The CMS Innovation Center's Strategy to Support High-quality Primary Care (includes 2000–2019 multi-physician exposure stats).

16 Arndt BG, Beasley JW, Watkinson MD, et al. Tethered to the EHR: Primary Care Physician Workload Assessment Using EHR Event Log Data and Time-Motion Observations. Annals of Family Medicine. 2017;15(5):419–426. HTML | PDF

17 American Medical Association (AMA). Physician burnout rate drops below 50% for first time in 4 years (reports 48.2% with ≥1 symptom in 2023). HTML

18 Credo Health point-of-care beta synthesis: 150+ clinicians; qualitative + product telemetry summaries. (Internal; not publicly cited.)

19 CDC / NCHS FastStats. Physician office visits (1.0B visits; 50.3% to primary care; cites NAMCS 2019). HTML

20 Dantas LF, Fleck JL, Cyrino Oliveira FL, Hamacher S. No-shows in appointment scheduling – a systematic literature review. Health Policy. 2018;122(4):412–421. HTML | DOI explicitly for the reference list