

# Why HL7v2 Still Wins for Pushing Data Into the EMR

**by Marilee Benson**

Modern health tech teams often build their first EMR integration on FHIR.

It's well-documented, uses familiar REST patterns, and every conference talk recommends it.

Then a customer asks them to push a result back into EHR, or update a patient's demographics after intake, and they hit a wall because the FHIR API is read-only.

This catches a lot of vendors by surprise, but it shouldn't. The path to writing data into an EMR today still primarily runs through HL7v2-type messages, and there are good reasons for that. The teams shipping integrations fastest are the ones who treat HL7v2 as a feature of the standards landscape, not a workaround.

## The read-only reality of FHIR

When ONC published the 21st Century Cures Act Final Rule in 2020, the FHIR API certification criterion had a specific goal. Patients needed secure access to their own health information through third-party apps. The rule [explicitly limited the new criterion to API-enabled "read" services](#) using FHIR Release 4.

HIMSS [confirmed the same point](#) in its analysis. The certification criterion required standardized API access "limited to API-enabled 'read' services." There was [no requirement for an API user to update data in the EHR](#).

EMR vendors built to the requirement. Most haven't invested heavily in bidirectional FHIR because the regulation didn't ask them to, and retrofitting a legacy system to write through FHIR is expensive. Epic and other EHRs like eClinicalWorks support some write capabilities for specific resources, but [resource support and read/write permissions vary by site and FHIR version](#). Confirming what each customer site actually allows is one of the most common scope-blockers in EHR integration projects.

There's a second wrinkle that catches teams off guard: **Cost**.

Read access through FHIR is typically free, because that's what the regulation requires. Write access is often a different story. Even when an EMR's FHIR or proprietary API supports “writes”, those capabilities often come with usage fees. A common structure is a monthly base plus per-transaction fees that vary by data type. And fees can be substantial with fees of +\$200 a month for relatively low message volumes for a single write type. Multiply that across data types and volumes, and the variable cost becomes a real factor in the architecture decision.

This is where FHIR sits in its maturity curve. It aligns with regulation very well. But writing into the EMR is lagging behind, and when that feature exists, the economics aren't always favorable.

## Where FHIR is still maturing

FHIR works well inside a single organization and within a vendor ecosystem. Epic [reported at HIMSS 2025](#) that FHIR API usage across its community surged from roughly 6 billion to over 10 billion monthly calls in a single year. That's a real proof point for what FHIR can do at scale.

The harder problem is cross-organization exchange. Several pieces are still being figured out:

- A national directory for discovering and connecting FHIR endpoints across organizations
- Robust “event” subscriptions for event-driven workflows between organizations
- Mature bulk FHIR implementations beyond a handful of leading EMRs
- Granular resource-level filtering for protected data categories like behavioral health, maternal health, and gender identity
- National networks / trusted participants security model

You can hear the uncertainty in the conversations happening right now. One stakeholder will describe a clear path forward using a particular standard or approach. Another, on a different call the same week, will say that's not finalized and that a different approach is emerging. Some industry voices argue that adopting FHIR too quickly could actually set the field back, because document-based exchange (CCDA, for example) gives clinicians a more complete view of a patient



than a series of targeted FHIR resource calls. Strip the document down to discrete resources, and you risk losing the one piece of context the clinician didn't know they needed.

None of this is a knock on FHIR. The standard is evolving, the ONC is working through what the next round of rulemaking looks like, and the industry is finding the right architecture for national-scale exchange. It just means CCDAs and HL7v2 aren't going anywhere soon, and the smart move is to support all three.

## What HL7v2 does well

HL7v2 has been driving data exchange in healthcare since 1989. According to HL7 International's own product brief, [95% of US healthcare organizations and providers in more than 35 countries](#) use HL7v2 as their workhorse for clinical data exchange. The standard is mature, actively maintained, and tightly embedded in how EMRs actually run.

Three properties make it the right tool for write-back use cases.

**It's bidirectional by default.** Almost every HL7v2 message type works in both directions. ADT, ORU, ORM, MDM, and SIU. You can send or receive, and the EMR knows what to do with it.

**It's event-driven.** When a patient is registered, an appointment is booked, an order is made, or a result is finalized, the EMR produces or receives a message. Your system doesn't have to poll. The trigger lives where the workflow lives.

**It's wired into the clinical workflow.** A HL7v2 ORU result lands in a flowsheet. An MDM document lands in the chart. An ORM order lands in the order queue. These aren't side channels. They're the same paths the EMR's internal systems use.

## The use cases that need HL7v2 today

A short list of what modern health tech vendors actually need to push into the EMR:

- Lab or diagnostic results from a clinical tool, sent as ORU
- Transcribed notes or AI-generated documentation, sent as MDM
- Updated demographics from a digital front door, sent as ADT
- Orders generated by a clinical decision support tool, sent as ORM



- Referral status updates that close the loop with the referring provider, via the 360X pattern using HL7v2 order messages and C-CDA documents over Direct Secure Messaging

For each of these, FHIR is either not supported for writes or requires per-site negotiation to enable. HL7v2 is the path with the fewest surprises across a wide range of EHRs.

## Standards work together, not in competition

The best integration patterns use each standard where it shines.

A common architecture looks like this. An ADT feed tells your system when something happened and who the patient is. You use the EMR's FHIR API to pull the granular clinical context you need. Or, you can trigger a national network query for the full patient record in a Treatment scenario. When you have something to write back, a result, a document, or an order, you push it via HL7v2. Direct secure messaging often sits alongside this for closing the loop with referring providers. [Industry guides describe the same hybrid pattern](#): FHIR for structured data queries, HL7v2 for real-time event-driven messaging.

This pattern works because it respects what each standard is good at. ADT is a reliable, low-cost notification layer. FHIR is a clean way to query specific resources once you know who you're asking about. HL7v2 is the path the EMR already trusts for inbound data. Direct messaging covers the workflow gaps that neither one solves cleanly.

Vendors who design around this hybrid model ship faster. They avoid the trap of building twice when a customer's FHIR endpoint doesn't support the write they assumed, or when transaction fees make the FHIR write path uneconomic at their volume.

## Making HL7v2 easier to live with

HL7v2 has real friction. Every EMR has its own dialect. Pipe-delimited parsing is unpleasant. Standing up a new connection often requires VPNs, license or setup fees, and back-and-forth with the customer's network engineering team. None of this is a reason to avoid HL7v2. It's a reason to abstract it.



The mistake we see often: a modern health tech team hard-codes HL7v2 support directly into their product. Every new customer becomes a project. Every field-level variance becomes a release. The better path for modern health tech teams is to use a clean, modern interface, JSON over a REST API, and let a bridging service handle the HL7v2 generation, per-site localization, and message transport.

That's why we built the [Zen HL7v2 <>JSON service](#). Modern health tech vendors connect with JSON. Zen's Gemini platform handles standards-compliant HL7v2 generation, per-customer localization, and transport. The integration team gets to focus on product workflows instead of field 19 of an OBX segment.

The point isn't that HL7v2 is hard. It's that you shouldn't have to rebuild a HL7v2 stack from scratch for every new hospital or provider you onboard.

## The takeaway

FHIR earned its reputation. It opened up patient access, supported a generation of consumer apps, and gave developers a modern API surface to work with. The standards landscape is healthier with FHIR in it, and the work being done now to scale it across organizations is important.

If your product needs to write to the EMR today, or be sent a message when an event occurs in the EHR, HL7v2 is still the answer for most use cases. Treat it as a tool, not a tax. Combine it with FHIR where it makes sense. And find a partner who can take the integration plumbing off your plate so you can ship new features

Want to learn more about how we help organizations like yours navigate the complexities of interoperability? Book a call with one of our experts by [clicking here](#).

