OpenEMR Weekly Meeting Detailed Summary

November 10, 2026

Attendees: Stephen Waite, Stephen Nielson, David Eschelbacher, Jerry Padgett, Simon Quigley, Prabhu Pathak (LogiCoy), Eric Stern (OCE), Seth C, Jit Chawla (ACE Health Solutions), Sherwin Gaddis, Farjad Zaghi (joined near the end).

Not attending: Brady Miller, Ansuman Nayak (LogiCoy), Michael Smith (OCE), Pedro Camellon.

> *Italicized indented passages are background, clarification, and commentary added after the meeting, not statements made on the call.*

- **Shape of the meeting:** Stephen Nielson was back after missing last week and chaired what he called a technical discussion. The call ran about an hour. It covered the architecture documents, welcomes for a new contributor and a returning one, Jerry's FHIR write and audit work, EHI export and patient merges, Direct messaging and referrals, Jerry's new telemetry controls, Docker versus multisite, and telehealth recording. With Brady absent, the PR review was short; Stephen hopes to have him explain his PR-shepherding approach next week.
- **SMART scopes, settled:** Stephen had replied on the forum to Jerry's question from last week about refusing SMART v1 scopes for new clients. Jerry is dropping it — it isn't worth the debate, since legacy clients have to be supported anyway. Stephen recalled the same debate a year ago: the project had to keep supporting v1, which is why it built the permission-mapping layer he is still not happy with. Jerry said his recent trouble came from other issues, not the scopes themselves.

> *This closes an item carried from 11/3. OpenEMR's API documentation labels v1 scopes "deprecated but supported" and maps them to v2 internally — .read to .rs and .write to .cud — which is the mapping layer Stephen meant. Keeping it means v1 clients continue to work, while new work, such as the granular category scopes needed for FHIR writes, can target v2.*

Architecture Documents — Nearly Ready

- **Stephen's goal:** make sure everyone is comfortable with the architecture documents. He believes he has addressed the concerns Jerry and Eric raised.
- **Jerry's review:** his re-review is on hold because he has since pulled all the Phreeze code out of the portal, so he will update his comments to drop those points and post them after the call. Eric, who had read about half the documents, asked whether Jerry's concern was about not changing too much. No, Jerry said: his concern was that they stay "too close to the surface" and don't go far enough down into the code — mostly in the portal, which the Phreeze removal has largely addressed. Eric called that totally valid. Otherwise Jerry considers them good, and expects them to help a lot with AI.
- **A living document:** Stephen called this a first draft to keep revising, like AGENTS.md: when the calls make an architecture decision that corrects course, the documents get updated. He has been tied up with client work, and expects to get back to the documents tomorrow.
- **Eric's suggestion:** if one area is uncertain or contentious, leave the portal out of the initial version, land the rest, and cover the portal later. Jerry said there isn't much more to do in the portal anyway.
- **Outcome:** once Jerry does another pass, get it in, so AI tools can start using it. Stephen is already seeing discrepancies with the architecture in PRs he reviewed that morning: "the sooner we get it in, the bigger the benefit."

> *Context from 11/3: Jerry objected to a rule in the draft that, in effect, said a working dependency should not be removed, and the Phreeze removal he described then has now changed what the portal section needs to say. Eric's suggestion is a common way to keep one contested area from holding up the rest: merge what is agreed, and mark the portal section as in transition until its new shape settles.*
>
> *Two practices fit the "living document" goal. Architecture decision records — short dated notes giving the context, the decision, and its consequences, a format popularized by Michael Nygard in 2011 — would let each course correction from these calls be recorded once and linked from the main documents, so readers see both the current rule and why it changed. And because coding agents read files such as AGENTS.md literally, an out-of-date rule misleads an agent more confidently than a missing one; a habit of updating the documents in the same PR that changes a convention keeps them trustworthy.*

New and Returning Faces

Seth C — A New Contributor Who Builds with Claude

- **Welcome:** Stephen welcomed Seth to his first call. Seth confirmed he is anun333 and Stephen said it was good to put a name to a face: the maintainers have appreciated his work.
- **Background:** Seth is a maintenance worker at a healthcare company with a strong interest in healthcare, and no prior programming background or experience. He started by using Claude to create OpenEMR modules, and since early October has used Claude to make PRs of fixes and minor improvements to the core codebase. He uses Claude for all of his coding.
- **How he works:** he described his part as more like administrative work — making sure the AI covers everything it needs to look at and gets a full picture. He was glad OpenEMR accepts AI-assisted submissions when many projects don't.
- **Stephen on velocity:** the struggle is keeping up with the volume of submissions, which is why many projects are rejecting them. Outside contributions used to run three to five on a good week; this is the busiest the project has been since version 5. Contributors who respond to feedback are very welcome; drive-by AI submissions are not, and neither are PRs that slip things in — he closed one that day that advertised the submitter's product in the README.

> *Other projects have responded to the same pressure with new GitHub controls. In February 2026 GitHub added repository settings to turn pull requests off entirely or limit them to collaborators, and in June it let maintainers cap the number of open PRs from users without write access. PRs opened by AI agents count toward the cap, and a bypass list exempts trusted contributors. GitHub's stated rationale: "The cost to create has dropped but the cost to review has not." A per-contributor cap with a bypass list is the closest match to Stephen's position: it keeps the door open while limiting how much any one newcomer can add to the queue at once.*
>
> *Seth has appeared in these notes before as anun333: PR 15012, which fixed creating insurers through the FHIR API (11/3); the FHIR Task resource, which Jerry credited to him; independent verification of Farjad's PR 14007 (10/27); and review comments on Farjad's DICOM PR this week.*

Jit Chawla Returns

- **Who he is:** Jit runs ACE Health Solutions in the Raleigh-Durham area of North Carolina and was an OpenEMR support provider for many years. Prior to his recent one-year hiatus, which began with his formal resignation from the board in July 2025, he served for years as the Vice President of the OpenEMR Foundation Board of Directors.
- **His request:** he is reviving his company and wants it back on the Professional Support page at Tier 3. He made a donation through the Foundation's donation service at the start of the month and added his website, and had emailed Brady and Stephen Waite.
- **Stephen Waite:** he remembered the email. Jit just needs to get back into the wiki and edit his entry; Stephen will email him the instructions rather than take time on the call.
- **Welcome back:** several people welcomed him. Jit has been brushing up on AI and has ideas for how to incorporate it into OpenEMR, which he expects the group is already discussing. Later he asked about telehealth (below).

> *The Professional Support page lists organizations that make ongoing contributions to the project, in three tiers by monthly donation: Tier 1 at $100 or more, Tier 2 at $50, and Tier 3 at $20. GitHub Sponsors is named as the preferred method, and instructions for placing an entry are on the page's Talk page. Neither ACE nor Jit is currently listed. ACE's website describes OpenEMR implementation consulting: installation and cloud setup, data migration, custom modules, role-based staff training, and workflow optimization.*

Jerry's FHIR Work — Audits, Assessments, and What's Next

Merged This Week

- **Write audits are in:** Brady merged Jerry's patient FHIR write audit and the confidential flag the night before, so there is now a flag for both.
- **A new home for import:** Jerry moved import out of the portal dashboard and into the assessment dashboard, together with the document template — the same rendering as the portal — and availability, which lets you choose which document to present and where, with categories. It is all under Assessment, and he encouraged people to try it and help debug it.
- **Visuals:** Jerry put a few screenshots in one PR, showing how a write is approved or not and how it is presented. He plans a forum thread after the call with instructions, and hopes community testers will go through it as they usually do.
- **Stephen:** screenshots and documentation will be wanted eventually. He wants to look through the audit side, which he and Jerry have discussed at length. Jerry: it isn't set in stone, and he built it in phases so that backing it out isn't difficult.

> *This is phase 2 of the three-phase audit epic Jerry described on 10/6 — hold-and-review for writes from patient-facing SMART apps — with phase 1, the portal audit rebuild, as its groundwork (11/3). Building it in phases behind a flag is the safe way to put a new control on patient-facing writes into core: sites can turn it off, and a problem found in the field doesn't require a revert.*
>
> *On the confidential flag: in FHIR, sensitivity is conventionally carried in a resource's security labels, using HL7's confidentiality codes (N normal, R restricted, V very restricted). Whether the new flag maps to those labels in FHIR responses is worth checking in the PR, since that is how a receiving app would know to treat the data differently.*

EHI Export and Patient Merges

- **Keeping EHI export current:** Stephen asked whether the EHI export is being kept up to date as columns change. Jerry said yes: his last set of PRs included fixes for changes made by him and others, and he tests it constantly in his API work.
- **What about merges?** Stephen wondered whether the new audit pieces account for patient merges and duplicates; Seth had found a bug in the merge in a recent PR. Stephen floated re-architecting the merge to crawl the EHI export's map of tables, and said he might be able to find sponsorship for that work.
- **Seth's view:** he lets the machine do the work and is still learning a giant project. He would like to continue the modularization push: with things more isolated, upgrading to new conformance specifications should be easier. Stephen: OpenEMR followed “the big ball of spaghetti design pattern” for a while, and the project is now consolidating many different architectures and designs.

> *Electronic health information export, criterion (b)(10), has been required of certified health IT since December 31, 2023, for both single-patient and whole-population exports, and its format must be publicly documented. HTI-5 left it unchanged: it appears in neither the removals nor the revisions. That makes the export's table inventory one of the few lists in OpenEMR that certification forces to stay complete.*
>
> *Why that matters for merging: OpenEMR's merge (Admin > Patients > Manage Duplicates, with “Merge and Discard” and “Merge with Encounter De-Duplication”) moves records from one patient to another and is non-recoverable, so the wiki advises a backup first. The classic failure is a table added after the merge code was written that still points at the discarded patient. Driving the merge from the same inventory the EHI export maintains would close that gap. One caution: audit records should not be rewritten by a merge. The Security Rule's audit-control standard and the certification criterion for auditable events and tamper resistance both assume the log is a faithful history, so a merge should add a merge event linking the two records rather than altering earlier entries — including any FHIR writes still waiting in Jerry's review queue.*

Which FHIR Writes Next?

- **Jerry's next resources:** vitals are done, and a draft PR for Observation labs is up. He is debating whether to do surgery, meaning outpatient procedures. Stephen saw no reason not to, but said to work on what interests him.
- **Procedure writes:** Jerry wasn't sure why Procedure doesn't allow FHIR writes; it was a decision made a couple of years ago that he doesn't think holds up now. But orders go through ServiceRequest anyway, so he isn't sure it should be a focus.
- **Stephen's priority:** parity with the reads for everything US Core-compliant on 8.0. Beyond that, there are many implementation guides. A FHIR e-prescribing workflow would be valuable, and that would probably need the subscriptions and webhooks guide; nobody has submitted a PR for subscriptions yet.
- **Pulling things out:** Stephen has lately been focused on removing code from core. Pulling out the telehealth piece is on his list, and there are other candidates. Jerry needs a few core events for his fax module (heard as “facts”); once he has them, its dependencies are essentially done.

> *On US Core versions: Stephen's “8.0” most likely means US Core 8.0.0. HTI-1 made US Core 6.1.0, with USCDI v3, the certification baseline from January 1, 2026. ONC's 2025 Standards Version Advancement Process approved US Core 8.0.0 and USCDI v5 for voluntary use from August 29, 2025. US Core mostly defines reads and searches, so “parity with reads” is a project goal rather than a certification requirement: every resource OpenEMR can serve becomes one it can also accept, where that makes clinical sense.*
>
> *Jerry's Procedure point matches US Core's split: ServiceRequest is an order, Procedure is something performed. The case for Procedure writes is mostly outside history — a surgery done elsewhere and reported back by a patient app or another system, which today has to be keyed in by hand.*
>
> *For subscriptions, the HL7 Subscriptions R5 Backport guide (STU 1.1.0) lets FHIR R4 servers implement R5-style topic-based subscriptions, including a REST-hook channel — in effect, a webhook when a matching resource changes. Certified e-prescribing still runs on NCPDP SCRIPT (version 2023011 by December 31, 2027, under HTI-4), so FHIR subscriptions would complement rather than replace it, for example by notifying an e-prescribing module when a patient's medications or allergies change.*

Direct Messaging and Referrals

- **Stephen's question:** since the certification requirement for Direct is going away, and Michael has been refactoring the Direct code into classes — Stephen wasn't sure whether that was an OCE internal initiative — should Direct move into a module instead, so sites that don't want it can leave it out? He saw it as a stepwise process.
- **Jerry:** Direct is a pretty big part of his referral work. The recent Direct fixes arrived just in time — referral was one of his action items, and Direct is how tasks communicate between providers, for reporting and for generating C-CDA documents. There are other ways to do it, but Direct was already there, so he used it. He didn't think moving it into a module would hurt.
- **The referral workflow:** the referral form automatically creates a task, if that is turned on in globals, and it works end to end. Brady pulled in the PRs for the outside part of the referral — where Direct comes in — the night before. They aren't big PRs, but Jerry called them important and suggested Stephen look at them.
- **Transport-agnostic:** Stephen recalled TEFCA-related efforts to couple UDAP with Direct messaging for referral processing. From an international perspective, he said, the transport should be agnostic. Jerry agreed: Direct was simply the one at his fingertips.

> *A correction on timing: HTI-5 removes the stand-alone Direct criteria, (h)(1) and (h)(2), effective December 28, 2026. But the transitions-of-care criterion, (b)(1), keeps its C-CDA and Direct requirements until January 1, 2030, and OpenEMR's certification lists EMR Direct as relied-upon software for (b)(1) (11/3). So Direct remains a certification item for three more years. Moving it into a module still fits, as long as the certified bundle includes it — the same pattern the module policy (10/13) set for vendor-specific code.*
>
> *Jerry's design resembles ONC's 360X closed-loop referral project, now IHE-approved profiles: a referral gets a unique identifier that persists until the loop is closed, documents travel as C-CDA over Direct, and status updates use HL7 v2 messages. A transport-agnostic version keeps that state machine — task, status, documents — separate from the delivery channel, so Direct, FHIR (HL7's FHIR referral guides such as BSeR, built on ServiceRequest and Task), fax, or a national network outside the US can be plugged in. TEFCA's FHIR-based exchange relies on UDAP for app registration and trust, which is the coupling Stephen remembered.*

Telemetry — Choosing What Counts as Raw

- **Jerry's change:** he added a panel to the telemetry app that lets you decide what counts as raw data. About ten areas can be parsed out of the report. For geography, for example, you can publish only the region, region and state, or everything. It is designed to meet people's apprehension about serving raw data publicly, which is where things were left last week: “pretty much any way you want to give that data out, you can do it.”
- **Response:** one participant voiced enthusiasm for making it available — it has been discussed for a long time, and they were happy to see movement.

> *This is the technical half of last week's question. The other half is a policy: which level is published by default, to whom, and how small counts are handled. Choosing fields is not quite the same as anonymizing. Even at state level, a state with one installation reporting a rarely used feature identifies that installation, which is why public health releases commonly suppress any cell below 11. A minimum-count threshold in Jerry's panel, applied before anything leaves it, would make the coarser settings genuinely safe.*

Docker, Multisite, and Modules

- **Docker over multisite:** one participant, following up on a question Stephen had asked on GitHub about modules, said that as the project pushes people toward Docker, most of the need for multisite goes away: you just spin up another container.
- **Shared hosting:** multisite mainly helps people on shared VPSs, a model the speaker didn't want to support much longer — the risks are too high. The recent PHP emergency showed why: on shared hosting you wait for the provider to update PHP or a vulnerable library, while a new image can be built, shipped, and applied within hours.
- **Compliance:** the view was that you can't run a compliant production deployment of OpenEMR on a shared VPS; an audit would find all kinds of violations. Running it that way to experiment is fine — nobody will stop you — but not with real PHI.
- **Stephen's caution:** the project still lacks a good answer when it tells people to drop multisite for Docker. Many will ask how to get their modules working and their custom configuration set up. That is easy for anyone who knows Docker and Docker Compose, but hard for traditional bare-metal administrators. He is seeing that now with a client's system administrators. The community needs clear instructions and training.
- **Revisiting the Docker tooling:** one developer had made a first pass, six or more months ago, at streamlining the Docker tooling and the development and production images. It was paused because it was blocked by the Node C-CDA service. Now that OpenEMR is pure PHP, it can be revisited.

> *Context: multisite and subdirectory deprecation has been carried on these agendas for several weeks, and the Node C-CDA service left the runtime with 8.5.0 (10/27), which is what unblocks the image work. The missing instructions Stephen describes are short. The usual pattern is a small derived image — a Dockerfile starting FROM the official openemr/openemr image, then a composer require for a module published on Packagist (the module policy of 10/13) or a COPY of a custom module — with site configuration and documents kept on volumes and the image rebuilt for each release. The common trap is installing modules into a running container: anything not on a volume disappears when the container is recreated.*
>
> *On compliance, the issue is the host's role rather than the word VPS. HHS's 2016 cloud computing guidance says a cloud provider that stores ePHI is a business associate, even if it holds only encrypted data without the key, so the host must sign a business associate agreement; most low-cost shared hosts will not. Add the Security Rule's risk analysis and timely patching, which shared hosting makes hard to control. A dedicated server or VPS from a provider that signs a business associate agreement can be compliant; a shared one usually cannot.*

Telehealth — Comlink, Custom Builds, and Recording

- **Jit's question:** what is the latest telemedicine offering that integrates with OpenEMR? Sherwin: it is still Comlink, and nobody else has posted a public one. Another participant has been working on a telehealth module but isn't ready to release it.
- **Stephen:** Sherwin has been the main support provider for Comlink, and it is unclear how heavily it is used. Many people have built their own Jitsi integrations; forum posts mention them, but say “contact us for details” rather than open-sourcing them.
- **Rolling your own:** one view was that many people now go their own way, since AI makes it easy to spin up a video integration — and when someone builds one, others ask whether they're giving it away. Another participant who runs an open-source project said that as soon as it went up, it drew drive-by AI contributions out of line with the project, and they had to turn off public pull requests.
- **What belongs in core:** telehealth is a large part of modern healthcare, so there should be some built-in support, as with payments, with specific vendors brought in as modules. People may keep private modules and duplicate each other's work, which the project can't prevent, but a released module like Comlink serves as a guide. The key is a stable module API, so that private modules don't block point releases and the security fixes that keep coming.
- **Where should a recording live?** one participant building a module on 8x8's hosted Jitsi service has added recording. When a session is recorded, the video sits with 8x8 and isn't pulled into OpenEMR automatically. Should it stay there, given HIPAA? 8x8 does offer a business associate agreement.
- **Stephen's answer:** people have to be careful rolling their own. A video service that only passes the stream through avoids business-associate obligations; once anything is recorded, you need business associate agreements and must meet all the privacy and security rules. Comlink specifically does not record — that was worked through during its security audits — unless a customer enters an enterprise agreement. Vendors generally charge more for recording because of the security and liability, and a recording raises questions such as patient requests and the designated record set.
- **Summaries too:** recordings invite AI-generated visit summaries — the way these call summaries are produced from a transcript — but that means choosing which partner does it. Someone remarked that the project really needs a lawyer on these calls.

> *HHS's 2022 guidance on audio-only telehealth draws the line Stephen describes. A service with only transient access to the PHI it transmits — one that connects the session and does not create, receive, or maintain PHI — needs no business associate agreement. One that stores PHI does. Technologies that electronically record or transcribe a session fall under the Security Rule, and OCR asks covered entities to assess whether stored recordings are encrypted. The pandemic-era enforcement discretion that let providers use any non-public-facing video tool ended with the public health emergency in May 2023.*
>
> *8x8's developer documentation answers the practical question. 8x8 states that it provides business associate agreements. It keeps a recording in its cloud for only 24 hours, then deletes it permanently, and it announces each finished recording with a webhook carrying a pre-authenticated download link. So a module can't leave recordings with 8x8: it has to fetch each one within the day — most naturally into the patient's documents in OpenEMR — and at that point OpenEMR is holding the recording. That brings encryption at rest, a retention policy, and the patient's right of access if the recording is part of the designated record set (records used to make decisions about the patient). Recording consent matters too: roughly a dozen states require every party to agree to a recording. Comlink's choice not to record avoids all of this.*

PR Review

- **Stephen's review queue:** he apologized for holding up PRs and is reviewing with the time he has. He looks only at PRs that are not in draft and have the green checkmark, or are close — 80 or 81 of 82 checks. If one or two checks are off, he reruns the end-to-end tests, which sometimes fail.

PR 8865 — Patient Finder Update

Pull request: https://github.com/openemr/openemr/pull/8865 (by psoas — David)

- **Background:** Stephen's own patient finder refactor, which he believed built on David's earlier work, was merged recently. He needed it for a client deployment that requires additional national identifiers, with encryption rules governing what can be displayed and who can see it. It added combined columns, row details, and column widths to the patient finder.
- **Still open:** Stephen had thought 8865 was merged; it dates from about a year ago. David explained that the refactor onto current code is done and CI is passing, though some UI styling still has to be reintroduced. What remained was resolving eight old review comments from Copilot that refer to code no longer in the PR, which he could not resolve himself. He also noted that contributors who aren't maintainers can't tick the checkboxes CodeRabbit posts to trigger a review it skipped. The comments were resolved during the call, so David can finish adding back the remaining features — the UI styling and making sure the columns resize correctly.
- **Aside:** Stephen has opted into GitHub's new PR review interface — actually two of them — and isn't sure he likes it.

> *From the PR thread: 8865 "adds many features to patient finders, including increased customizability" — expandable layout options, an exact-match search toggle, new columns such as address, appointments, age, and insurance, and an updated color palette — and fixes issue 8864. David first submitted it in 2022 and later reopened it. Brady's AI-assisted review lists blockers and estimates one and a half to two weeks of focused work to make it mergeable, and Michael noted that the stylesheet's heavy use of !important is a problem across the repository rather than one this PR can fix. David's latest comment raises a concern worth carrying forward: feature PRs from outside contributors tend to stall.*
>
> *Some mechanics that may help. GitHub lets the PR's author, or anyone with write access, resolve review conversations, so outdated bot threads on one's own PR should be resolvable by the author; if the button doesn't appear, a maintainer can do it. CodeRabbit's checkboxes live inside its own comment, which only people with edit rights can change, but anyone can type its commands instead — "@coderabbitai review" or "@coderabbitai full review." Replying "@coderabbitai resolve" inside one of its threads resolves that thread. Given the size of the remaining list, splitting 8865 into smaller PRs — the new columns first, then the layout and theme — would also fit the green-checkmark filter Stephen described.*

Flaky End-to-End Tests

- **The question:** a participant asked whether there is any way to fix the flaky tests, which have been a pain for as long as they had contributed. The project has gone back and forth on the end-to-end tests: whether to drop some, or switch to a different format.
- **Playwright or PHP?** Stephen: they are native PHP. He has tried both Playwright and PHP tools on internal projects, and both were flaky — anything that waits on requests to finish. The real cause is the way OpenEMR handles sessions, cookies, and tabs, and until that is refactored the tests will stay flaky. Another participant: "the entire auth system."
- **CI scope:** one test is usually the one that fails. A participant believed PRs run the end-to-end tests on a single PHP version, with the other versions only on master, to avoid swamping the runners.

> *OpenEMR's end-to-end suite lives in tests/Tests/E2e and uses Symfony Panther, a PHP library that drives a real Chrome browser, as its phpunit configuration shows. Stephen's diagnosis fits common experience: Playwright's automatic waiting removes timing flakiness, but not flakiness that comes from shared state in the application itself, such as sessions and frames that outlive one test. Usual mitigations are waiting on explicit conditions rather than fixed delays, a fresh login and session for each test, and a quarantine list for known-flaky tests that tracks them rather than silently retrying. Logic that doesn't need a browser can move to the faster API and web UI suites.*

PR 14254 — Sending NPIs as Digits

Pull request: https://github.com/openemr/openemr/pull/14254 (by tsimonq2 — Simon)

- **Choosing a PR:** Stephen asked Simon or Seth for a PR to discuss. Simon has 11 open. Someone asked whether the regional concern about NPIs had been resolved; it was clarified that 14254 is the billing fix, not the separate PR that would require an NPI.
- **Stephen Waite:** asked to drive it or comment on the 837 changes, he said he would need to look at it. He recalled that last week Brady simply screen-shared and stepped into his Claude console, which he thought was the impetus for Brady's shepherding work. Time ran out before a review.

> *The PR, opened September 25, fixes 837 claims that sent an NPI with a stored hyphen: the X12 cleaning function keeps hyphens, which suits names and streets but not NPIs. Its scope has grown to blank NPIs with invalid characters, apply the same rules to UB-04 physician fields, omit zero NPIs, and change five API test clients — a change reviewers may want separated. Stephen Nielson asked Stephen Waite to review it on October 2, and it has been on these agendas since 10/20. An NPI is ten digits, the last a check digit computed with the Luhn formula over the prefix 80840, so validating the check digit when an NPI is entered would stop bad values from being stored in the first place. A blanked NPI will be rejected by the clearinghouse, which is safer than sending a malformed one, but the biller should see why.*

Brady's PR Shepherding and the AI Review Loop

- **Stephen Nielson's question:** having missed last week, he asked whether the group had found it more productive to follow Brady's workflow than to review PRs by hand. The explanation offered: Brady is trying to find the best way to say a PR is good to go or not, balancing Claude and CodeRabbit. Depending on the PR, he may run a Claude review after CodeRabbit's, then CodeRabbit comments again, and the back and forth has to stop somewhere.
- **The loop problem:** one developer called adversarial AI review a problem many people are struggling with. The night before, he had told Fable to build something he had discussed with it for an hour, and went away from the keyboard; it ran an adversarial review at every step and built completely the wrong thing.
- **Brady's approach:** one participant thought Brady is onto something with his shepherding. His PR is a work in progress, with Claude reviewing deep in the weeds; it should make it possible to search PRs by author or category and make sense of the 300-plus open PRs, about 100 of them recent. Stephen wants Brady on next week's call to explain the tools and approach, because it helps the community to know what the maintainers look at.
- **A GitHub front end:** one participant has thought about having Claude build a GitHub front end like the one Eric described. Eric's is working locally and he offered a place in the beta, but it is Mac-native, and that participant doesn't use a Mac.

> *The convergence problem has known remedies. An AI reviewer will always find something more, so the stopping rule has to come from people: fix correctness, security, and data-integrity findings in the PR, and file the rest as follow-up issues; set CodeRabbit's review profile to its quieter "chill" setting rather than "assertive"; cap the number of review rounds; and leave the final call to a human. Fable is one of Anthropic's Claude models. The story also illustrates Michael's point from 11/3: an agent left to make architectural decisions on its own tends to choose wrong, however much it reviews itself.*

PR 15070 — Farjad's DICOM Viewer Fix

Pull request: https://github.com/openemr/openemr/pull/15070 (by doctorman — Farjad)

- **Farjad:** he joined as the call was ending, mainly to see what had moved forward, and thanked everyone for the reviews. About 26 of his PRs have been merged. The document page "was horrible" and bothered him, and the DICOM viewer is now very usable on his site: he can page through x-rays and add them, and zip uploads that used to fail now open.
- **Screenshots:** Stephen had commented on the DICOM PR that morning, and asked that future UX PRs include screenshots; Farjad will add pictures.
- **Jerry:** the DICOM project is old, out of date, and no longer supported. It breaks about every six months as people drive by, and needs to be rewritten and pulled out into a module. He had held off working on it until a rewrite was decided.
- **Status:** Brady has merged most of Farjad's PRs, so the DICOM PR is the only one still open. Brady's shepherding work the night before was said to be already paying dividends.

> *The PR, opened November 5, makes a zip of DICOM images show one image at a time in the documents viewer, with previous and next buttons; before, only the first image loaded. It reads the zip in the browser with JSZip, skips folders and system files, sorts the images, and falls back to the old loader if anything fails. Seth reviewed it for rapid clicks, non-image entries, and ordering. Stephen requested changes on November 10, asking for tests; Farjad added 14 Jest tests, the checks are green, and he has asked for another review. Related work includes server-side zip handling in PR 15224 and several other documents-viewer PRs.*
>
> *OpenEMR's DICOM viewer is Jerry's integration of the open-source DWV (DICOM Web Viewer) library, documented on the wiki from the 5.0.4 development cycle. Farjad's fixes and Jerry's call for a rewrite as a module aren't in conflict: the fixes help users now, and a module would let the viewer be replaced or upgraded without touching core, in line with the module policy.*

Next Week's Agenda and LogiCoy

- **Module distribution:** one participant would like to work out the module distribution patterns, since the project is pushing people to Docker, and offered to put some ideas in writing to discuss.
- **LogiCoy:** LogiCoy's e-prescribing was proposed for next week's technical discussion. Jerry noted that LogiCoy has only just begun testing with a couple of users, and its reporting is based on their feedback, so it may take a week or two to gather enough data.
- **Fairness:** the group wants to stay fair to every e-prescribing vendor while taking advantage of LogiCoy's interest.
- **Brady:** on PR shepherding, if he can attend.

> *Context: on 10/27 the call described two paths for e-prescribing vendors — publish a community module, or pursue certified bundling through the Foundation — and on 11/3 LogiCoy reported pilots with a couple of vendors. Certified e-prescribing, (b)(3), must move to NCPDP SCRIPT 2023011 with real-time prescription benefit by December 31, 2027, under HTI-4. Written, published criteria for certified bundling would let the project work closely with one vendor while keeping the same door open to others.*

Closing

- **Stephen's thanks:** open source work can feel thankless, and people wonder whether anyone benefits from their volunteer labor. He thanked everyone for their contributions and what they mean to the community: the project doesn't always make the best decisions, but everybody is trying. The reply: maintaining OpenEMR isn't easy, but it depends on contributions — a two-way street.

Items Discussed and Proposed

Recorded as topics raised and directions favored, not as assigned work or commitments.

Area	Discussed / Proposed
SMART v1 scopes	Question dropped; v1 remains supported for legacy clients through the mapping layer.
Architecture docs	Close to final; to be merged after Jerry's final pass, then revised as a living document. Portal section could follow later.
AI contributions	Welcome from contributors who respond to feedback; drive-by and promotional PRs closed.
FHIR write audits	Write audit and confidential flag merged; import, templates, and availability moved under Assessment. Testing and screenshots wanted.
EHI and merges	EHI export kept current. Question raised whether merges and audit logs handle duplicates; idea of driving merge from the EHI table map.
FHIR writes	Observation labs in draft. Procedure and surgery writes undecided. Suggested priority: parity with US Core 8.0 reads.
Direct	Central to Jerry's referral workflow; module move possible later. A transport-agnostic design favored.
Telemetry	New panel controls what is published as raw data, down to geographic level.
Docker	Favored over multisite; clearer instructions needed for modules and custom configuration. Image tooling to be revisited.
Telehealth	Comlink remains the public option. Recording brings business-associate and privacy obligations.
Support listing	Jit seeking a Tier 3 listing; wiki instructions to follow by email.
PR 8865	Old review comments resolved on the call; remaining features to be added back.
PR 14254	Not reviewed; awaiting billing review.
PR 15070	Tests requested; since added. Screenshots requested for UX PRs.

Future Discussion and Next Meeting Agenda

Proposed for upcoming meetings, drawn from this week's discussion, chat, and notes; not a fixed agenda.

Item	Details
Module distribution	How modules and custom configuration work in Docker deployments; written ideas, if ready.
PR shepherding	Brady invited to explain his approach and tools; where AI review should stop.
LogiCoy	Pilot feedback, and a possible role in (b)(3) certification.
Architecture docs	Final pass and merge; how to keep them current.
Telemetry	What to publish by default now that the panel exists; small-count handling.
Merges and EHI	Whether merge should use the EHI table map; how audit logs treat merged patients.
FHIR roadmap	US Core 8.0 write parity; Procedure writes; subscriptions for e-prescribing.
Direct and referrals	Transport-agnostic referrals; when Direct might move to a module.
Flaky tests	Options for the end-to-end suite.
Telehealth	Guidance for module authors on recording and HIPAA.
PRs carried	8865, 14254, and 15070.
Carried over	HTI-5 cleanup plan; HTI-4 deadlines; core-module plan; CodeRabbit resolution rule; license-header CI; contributor priorities list.

Notes

- **The connecting thread:** making OpenEMR legible — to people and to AI. The architecture documents, the push for modules, a merge driven by the EHI table map, transport-agnostic referrals, Docker instructions for modules, and screenshots on UX PRs all aim at the same thing: a codebase whose shape is written down, so that contributors and their agents can work in it without guessing.
- **The AI-assisted contributor model is working:** Seth, with no programming background, and Farjad, a physician with about 26 merged PRs, are both producing useful work — because they respond to feedback. The bottleneck remains review, which is what Brady's shepherding is meant to address.
- **Brady's absence was felt:** the PR review was brief, and the explanation of his shepherding approach moved to next week.
- Ways to help: try Jerry's FHIR write audits under Assessment and report what you find; review open PRs that have passing checks; and if you run OpenEMR in Docker with modules or custom configuration, share how you set it up — it is exactly the documentation the project needs. Comments on the forum or at the weekly call are welcome.