

OpenEMR Weekly Meeting Detailed Summary

October 13, 2026

Attendees: Stephen Waite, Stephen Nielson (departed early), Brady Miller, Eric Stern (OCE), David Eschelbacher, Simon Quigley, Sherwin Gaddis, Ansuman Nayak (LogiCoy), Luis Santos, Prabhu Pathak (LogiCoy), Jerry Padgett. A Read AI notetaker bot was also present.

Italicized indented passages are background, clarification, and commentary added after the meeting, not statements made on the call. This week's summary carries more of them than usual, because the e-prescribing discussion turned on certification and regulatory context that is worth spelling out.

- **Shape of the meeting:** the largest attendance in a long time. The first half was given to an introduction and demo from LogiCoy, an e-prescribing vendor that has built an OpenEMR module. The second half was a PR review centered on the LBF statements module, which reached a decision. Items carried over from last week for Stephen Nielson's return — the multisite and subdirectory deprecation plan, and procedure writes over FHIR — were not reached.

Opening — An AI Notetaker on the Call

Link: [Read AI](#)

- **An unexpected attendee:** a “Read AI meeting notes” bot appeared in the participant list. None of the regulars had added it, and the new LogiCoy attendees said it was not theirs.
- **Traced to a participant's colleague,** who had been using Read AI for note-taking. Because the colleague's account is tied to that participant's meetings, the bot keeps following them into calls — “that thing is freaking sticky.” The participant apologized and offered to remove it, leaving the choice to the group, which did not pursue it.
- **The participant's assessment of the tool:** it helps to an extent. It picks up clearly stated action items and terms you need to remember, but its reports can paint a less nuanced picture than what was actually said, so they need editing before being relied on.
- **David Eschelbacher reminded the many new attendees** that the meeting is recorded and that he uses AI to produce the summaries posted on the wiki, offered to post the link, and asked to learn more about Read AI. The participant offered to send information.

Read AI is a Seattle company founded in 2021 by David Shim, Robert Williams, and Elliott Waldron. Its notetaker joins video calls as a visible participant, records and transcribes, and sends a summary report afterward. Its own help pages explain the “sticky” behavior: during onboarding a user chooses whether the bot auto-joins every meeting on their calendar, only meetings they host, or only meetings they invite it to. People who receive a shared report and click through to view it can end up with an account of their own, often with their calendar connected. The fixes are to change the auto-join preference in account settings, disconnect the calendar, or delete the account. On the host side, a bot can be removed from a Google Meet call like any other participant.

The participant's caution — useful for action items, superficial on nuance — applies equally to these wiki summaries, which are also AI-assisted. The difference is process: these summaries are reviewed before posting, and each one invites corrections. A third-party notetaker's report goes, unreviewed, to whoever the bot's owner chooses. On an open community call that is not a problem in itself, but a sentence on the call details page noting that third-party recorders and notetakers may be present, and that the wiki summary is the reviewed record, might avoid confusion.

Agenda and a Standing PR Slot

- **Simon Quigley has a good number of PRs open** and asked whether Brady wanted to follow up on the LBF statements module (PR 13795), first raised a month or two ago. Brady had not looked at it again, having been concentrating on security work, and suggested reviewing it during the call.
- **Two of Simon's PRs have already been merged.**
- **PRs in the middle of the call, not at the end:** with so many PRs arriving, many of them AI-assisted, the proposal was to take pressing issues first and then turn to PRs around 2:30–2:40 Eastern — perhaps half the call — with the goal of getting through one or two each week to keep things moving. Luis Santos also had a PR question.

This is a meaningful process change. On 9/15 it was said candidly that review had been triaged to bug fixes, with feature PRs left sitting, and the suggestion then was for someone to take on straightforward feature PRs. A standing slot in the weekly call is a structural answer: at least one or two contributor PRs get live attention each week, and contributors have a predictable place to bring them. Today's slot produced a decision on a PR that had been pending since mid-September.

LogiCoy — A New E-Prescribing Option

Links: [LogiCoy](#) · [LogiCoy eRx](#) · [OpenEMR ePrescribe wiki page](#)

Who LogiCoy Is

- **Introductions:** Prabhu Pathak handles security operations at LogiCoy, and Ansuman Nayak is its technical lead. LogiCoy is a California company with a development team in India.
- **Why they came:** Ansuman has built an e-prescribing module that can be plugged into any OpenEMR installation. LogiCoy wants to help the community with it, and to partner with OpenEMR to reach more users. They had already been in contact with Stephen Waite.
- **Direct Surescripts connection:** Stephen Nielson, introducing himself as OpenEMR's director of engineering, asked first whether LogiCoy connects to Surescripts directly or piggybacks on a reseller such as NewCrop. Answer: direct.
- **Two products:** a standalone, web-based e-prescribing application that a prescriber signs up for and uses directly; and an e-prescribing API that EHRs, EMRs, and telehealth platforms integrate. The same features are available through both.
- **Certification:** Stephen Nielson had looked up their website while the demo was being set up and noted ONC certification of the e-prescribing pieces through SLI. LogiCoy confirmed it holds the required certifications and has live clients on both products.
- **Open source pedigree:** Prabhu added that LogiCoy's founder and CEO, Prabhu himself, and another co-founder built the OpenESB product at Sun Microsystems, where the founder managed the OpenESB community — so they know how open source communities work.

LogiCoy was founded in 2009 by engineers who had developed OpenESB at Sun Microsystems. OpenESB is an open source enterprise service bus built on the Java Business Integration standard; after Oracle acquired Sun in 2010, an independent community formed to maintain it, and LogiCoy offered commercial development and support. The company has since moved into healthcare, with two distinct lines of business: e-prescribing, and state prescription drug monitoring programs.

LogiCoy's compliance page lists ONC certification (listing #10944 on ONC's Certified Health IT Product List); Surescripts certification for e-prescribing, medication history, and prescription benefit transactions; compliance with the DEA's rules for electronic prescribing of controlled substances (21 CFR Part 1311); and ID.me for identity proofing. SLI Compliance, which Stephen mentioned, is both an ONC-Authorized Testing Laboratory and an ONC-Authorized Certification Body — one of a handful of organizations that test and certify health IT under ONC's program — and LogiCoy also cites an independent security assessment by SLI.

A direct Surescripts connection matters because Surescripts operates the network that links nearly all US prescribers, pharmacies, and pharmacy benefit managers. Many e-prescribing features in smaller EHRs are white-labelled from a few certified intermediaries. A direct connection means LogiCoy holds the Surescripts certification and network agreements itself, which is what allows it to sponsor other software onto the network, as described below.

Demo — The Standalone Product

- Prabhu walked through the standalone web application, logged in as a prescriber.
- **Prescriber dashboard:** counts of patients and prescriptions, controlled versus non-controlled substances, and transaction types — new prescriptions, changes, refills, and cancellations.
- **Patients:** a recent-patients list, adding patients with demographics, and a prescription history for each patient.
- **PDMP built in:** the history view can include other states' prescription monitoring data, which various states require prescribers to review before prescribing opioids. LogiCoy is itself the PDMP vendor for Pennsylvania and Illinois, and now Vermont, with direct connections in about ten other states such as Kentucky. Where its competitor Bamboo Health holds the connection, the two collaborate to provide full coverage.
- **Writing a prescription:** choose the pharmacy, with its Surescripts service levels shown; search the drug database; set strength, supply units, and unit of measure; add notes; check electronic prior authorization; and add or confirm allergies and diagnoses. In API mode, patient context, allergies, and diagnoses would come from the EHR rather than being entered.
- **Safety checks:** drug–drug, drug–allergy, and drug–alcohol interaction checking, plus real-time prescription benefit lookup.
- **Controlled substances:** ID.me identity verification and two-factor authentication at signing. The API version works the same way, with the ID.me step handled through API calls instead of a screen.
- **Other features:** CancelRx and the other standard prescription flows; delegates, who can draft prescriptions for a prescriber to approve; reports; favorite drugs and order sets; user management; billing; and PDMP registration.

Surescripts service levels are flags in the Surescripts pharmacy directory that say which transactions each pharmacy can accept — new prescriptions, renewals, changes, cancellations, controlled substances, and so on. Showing them at pharmacy selection lets the prescriber see in advance whether, say, a pharmacy can receive an electronic controlled-substance prescription, rather than finding out from a rejection.

The DEA's rule for electronic prescribing of controlled substances (EPCS, 21 CFR Part 1311, in effect since 2010) sets four main requirements: the prescriber's identity must be proofed by an approved credential service provider; signing requires two-factor authentication; granting a prescriber EPCS permissions requires two people, one of them a DEA registrant; and the application itself must pass a third-party audit or certification. ID.me is one of the identity providers widely used for the first step. The rule also explains a detail from the demo: delegates may prepare a controlled-substance prescription, but only the prescriber can sign it.

Prescription drug monitoring programs (PDMPs) are state databases of dispensed controlled substances. Most states now require prescribers to consult them before prescribing opioids and certain other drugs, and a check from inside the prescribing workflow is far more likely to happen than one requiring a separate login. LogiCoy plays two roles here. It runs the state PDMP system itself in some states — Pennsylvania lists it as its PDMP system vendor, and Illinois moved its data collection to LogiCoy in June 2022 — and it sells integrations that let EHRs query PDMPs. Bamboo Health, the competitor mentioned, runs PMP Gateway, which gives EHRs a single connection to more than 40 state PDMPs.

Real-time prescription benefit (RTPB) shows the patient-specific out-of-pocket cost of a drug, and cheaper covered alternatives, at the moment of prescribing. Electronic prior authorization (ePA) lets the prescriber complete a payer's prior authorization inside the workflow rather than by fax and phone. Both are moving from optional to required, as described under the regulatory discussion below.

Inside OpenEMR — The Module So Far

- Ansuman then showed the module running inside OpenEMR. The basic prescribing flow is done, and the rest is being added.
- **Configuration:** the module appears in the module manager, with settings for the backend URL, token URL, a webhook URL for notifications coming back from the e-prescribing service, and credentials.
- **Prescriber registration:** a button registers a prescriber with the e-prescribing service directly from OpenEMR.
- **Prescribing:** drug search, pharmacy search, draft prescription, then send — the same flow as the standalone product.
- **ID.me and provider setup:** ID.me setup happens the first time a provider is configured, through the admin user screen. Ansuman added the fields the service requires to the add-provider form.
- **Notifications:** webhooks bring notifications back into OpenEMR — for example, when there is a problem with a prescription that the prescriber needs to address.
- **Still to come:** the advanced APIs, such as electronic prior authorization, which LogiCoy said it could build into the module in about a week.
- **Reaction:** Stephen Nielson called the integration really nice and was glad to see the webhook path in place.

Webhooks matter more in e-prescribing than they might appear. Much of the traffic is asynchronous and starts at the pharmacy end: change requests, renewal requests, responses to cancellations, and error messages all arrive after the prescription is sent. Without an inbound channel the EHR never learns of them, and the prescriber has to check the vendor's portal separately. The practical consequence for OpenEMR is that each installation needs an endpoint the vendor can reach, and incoming calls need to be authenticated — straightforward for hosted installations, less so for an on-premises server behind a firewall.

Two architectural points will matter in review. First, Eric Stern's order of preference from 9/29 — standard APIs first, then OpenEMR's own APIs, then a module, with core changes only to add a missing extension point — and his remark later in today's meeting that vendor-specific integrations should be their own Composer packages both apply here. Second, the extra provider fields: whether they are added through the module system's extension points or by modifying the core user screen will make a real difference to upgradability.

Certification — The (b)(3) Question

- **Stephen Nielson's key question:** has LogiCoy taken other EHRs through ONC's (b)(3) e-prescribing certification using its API? LogiCoy has about three API clients, two of them live, and obtained Surescripts certification for its API about two years ago.
- **Why OpenEMR has held back from bundling:** NewCrop and Weno integrations exist in OpenEMR, but OpenEMR is not currently certified to (b)(3). To certify, OpenEMR would have to declare that it uses LogiCoy's product as part of its certified configuration.
- **The Surescripts approval process, as LogiCoy described it:** Surescripts sends a questionnaire through LogiCoy to the EHR vendor; the vendor completes it; and Surescripts checks that the EHR is genuine and will not misuse the network — its concern is drug diversion — and approves it, typically in about two weeks. Integration then proceeds. LogiCoy expected OpenEMR, as a well-established EHR, to be approved easily.
- **A familiar model:** Stephen noted OpenEMR already works this way with EMR Direct for secure Direct messaging, and that OpenEMR was certified through Surescripts years ago, via NewCrop. He said the path described was further along than anything the project currently has for e-prescribing.
- **A note of caution:** one attendee doubted that certification would follow soon, and saw the integration as a value-add for now.
- **Stephen's position:** certification is not required to proceed, or even to bundle the module, but it is where OpenEMR would want to be, and he very much wants it. The cost of going through certification needs to be discussed.
- **Who handles Surescripts:** LogiCoy's US-based sales and enablement lead coordinates with Surescripts and would send the questionnaire.

“B3” on the call is ONC certification criterion 45 CFR 170.315(b)(3), Electronic prescribing. It requires a set of NCPDP SCRIPT transactions — new prescriptions, changes, cancellations, renewals, fill status, and medication history — along with supporting status and error messages.

The record bears out Stephen's account. OpenEMR 5.0 was certified as a Complete EHR under the 2014 Edition in February 2017, and its certificate listed NewCrop, EMR Direct's phiMail server, and two other products as additional software relied upon for select criteria. OpenEMR 8 was certified under the 2015 Edition Cures Update on January 30, 2026 (CHPL listing 11757). Its criteria do not include (b)(3), and its relied-upon software is EMR Direct, for transitions of care and Direct messaging. ONC's program lets a developer list such third-party software, which is exactly how a certified API like LogiCoy's would fit into OpenEMR's certification.

Timing matters. ONC's HTI-4 final rule, effective October 1, 2025, moves (b)(3) to NCPDP SCRIPT version 2023011, which brings electronic prior authorization transactions into the criterion, and creates a new criterion, (b)(4), for real-time prescription benefit. The older SCRIPT 2017071 can be used only through December 31, 2027. From January 1, 2028, any product certified to (b)(3) must use SCRIPT 2023011 and must also be certified to (b)(4). A certification effort starting now would be aiming at the new versions.

The extra work Stephen half-remembered later in the meeting is very likely safety-enhanced design, 170.315(g)(3). It applies to a list of criteria that includes (b)(3), and requires a user-centered design process and summative usability testing with real users, documented in a report submitted for certification. A separate criterion, (g)(5), asks developers to identify the accessibility standard they follow. Neither depends on the e-prescribing vendor; both would fall on OpenEMR.

Why certification matters to practices: clinicians reporting the Promoting Interoperability category of Medicare's Merit-based Incentive Payment System must meet an e-Prescribing measure — required, and worth 10 points — by transmitting prescriptions electronically using technology certified to (b)(3). Clinicians who write fewer than 100 permissible prescriptions in the period can claim an exclusion. Practices may assemble their certified technology from more than one product, so a separately certified e-prescribing product can fill the gap, but a single certified, integrated workflow is simpler for the practice, and is what Stephen is aiming for.

Pricing and Regulatory Readiness

- **Pricing:** LogiCoy presented its pricing on the call. It is per prescriber per month, with separate rates for prescribing controlled substances (EPCS), and different rates for the standalone product and for use through the API from within an EHR such as OpenEMR. LogiCoy's published standalone pricing is on its website. Commercial terms would form part of any partnership discussion.
- **Electronic prior authorization and new requirements:** asked whether LogiCoy is keeping up with the wave of new requirements coming from ASTP/ONC over the next year, LogiCoy said it is building them: Medicare Part D certification on what it called “version 60” by year-end, and an update to NCPDP SCRIPT 2023011. Surescripts sends notices with deadlines, and vendors that do not comply are shut off.

“Version 60” is the NCPDP Formulary and Benefit standard, version 60. In a June 2024 final rule, CMS updated the Medicare Part D e-prescribing standards: Formulary and Benefit version 60 and Real-Time Prescription Benefit version 13 are required from January 1, 2027, and NCPDP SCRIPT 2023011 — covering e-prescribing, medication history, and electronic prior authorization — from January 1, 2028. ONC's HTI-4 rule aligned certification with the same versions. LogiCoy's year-end target is therefore consistent with the first CMS deadline.

“ASPE” in the transcript is ASTP: in July 2024 ONC was reorganized as the Assistant Secretary for Technology Policy and Office of the National Coordinator for Health IT (ASTP/ONC). Both names remain in use.

The point about Surescripts shutting off non-compliant vendors is real, and is part of the value of a vendor relationship. The network, not the EHR, enforces version upgrades, and an intermediary that keeps current shields its EHR partners from much of that work. For an open source project with limited capacity, handing off that compliance treadmill is a significant benefit.

What a Partnership Would Look Like

- **Leadership discussion:** it was suggested that project leadership discuss the opportunity. Asked for the views of Brady and Stephen Waite, the response was positive — “sounds super promising,” with no reason seen not to press ahead — and it was noted that the project does not have many alternatives.
- **What bundling would mean:** Stephen Nielson explained that, as part of a certified product, the module would be in every OpenEMR installation, as NewCrop and Weno have been. He asked what LogiCoy has in mind for a partnership. LogiCoy will give a fuller answer on the next call, and is open to discussing what it can contribute to the community.
- **Not a vendor tier:** LogiCoy asked about the Tier 1, 2, and 3 listings on the OpenEMR website. Stephen explained that those are for vendors who install and support OpenEMR. LogiCoy would fit a different category of partner, like EMR Direct: call-outs on the certification pages, possible co-marketing, and, as the certified e-prescribing product, the main reference wherever OpenEMR discusses e-prescribing.

The tiers appear on the Professional Support page of the OpenEMR wiki, which lists vendors who provide installation, support, and customization. Placement is tied to a recurring monthly contribution to the project — Tier 1 at \$100 or more, Tier 2 at \$50, Tier 3 at \$20 — plus a set of listing criteria. A technology partner whose product is part of the certified configuration is a different relationship, and EMR Direct is the existing precedent: it is relied-upon software on OpenEMR's certificate, documented on the wiki, and integrated into OpenEMR since 2013.

Onboarding and Payments

- **The open question:** the module has nothing yet for onboarding — signing up and paying for the API service. Stephen Nielson asked how that would work, given OpenEMR's decentralized nature.
- **LogiCoy's current model:** API customers register in a vendor portal where they can see their transactions and prescriptions, and LogiCoy invoices them. If that does not fit OpenEMR, a billing module or payment integration could be built.
- **Likely shape:** installing the module would send the practice to LogiCoy's vendor portal to register and obtain credentials. LogiCoy offered a separate demo of the vendor portal.

Stephen's “decentralized” point is the crux of the partnership design. A commercial EHR vendor is a single customer of its e-prescribing provider: it holds the contract, onboards its practices, and bills them. OpenEMR has no central operator — thousands of practices install it independently, many through vendors — so it is not obvious who the “EHR vendor” is for Surescripts purposes, who completes onboarding, or who is billed. The existing options show two answers. Weno is set up directly by each practice, which pays Weno a sign-on fee and annual charges; according to the OpenEMR ePrescribe wiki page, part of the revenue from OpenEMR clients returns to the project. NewCrop, now Ensora Health, has been provided through a few authorized OpenEMR vendors, which hold the relationship and provide support. LogiCoy's vendor-portal model resembles the first, with each practice as a direct customer.

Next Steps

- Stephen Waite, who initiated the contact, will decide on a follow-up. Project leadership will discuss the opportunity without LogiCoy present and come back to LogiCoy afterward.
- LogiCoy proposed that the next meeting include its founder and CEO and its US sales lead, who works with Surescripts. Stephen Nielson said he would want to see a bigger demo later; LogiCoy noted the core APIs are in the module, with advanced APIs such as ePA still to come.

Reactions After LogiCoy Left

- **Broadly positive:** it looked official, they are certified, it is affordable, and they appear willing to get it going. It is the first option seen that is affordable and willing to work with the community; they did most of the work, and the integration so far looked decent.
- **A reservation:** one attendee said the Surescripts aspect bothered them, without elaborating.
- **Context from another attendee** who has been looking into e-prescribing recently, mostly for dental practices: many vendors simply connect to Surescripts and rent out a certified user interface, so it was good to see a real integration.
- **Certification work ahead:** Stephen Nielson recalled that certification previously meant updating a set of design and accessibility requirements, which would have been a major headache before AI. He was not sure they still apply, but they will need revisiting if the project pursues (b)(3). Overall, he thinks it is worth pursuing.
- **How LogiCoy found the project:** Stephen Waite explained that LogiCoy first posted on the forum. The forum's filters held the post — prepared text pasted in quickly gets flagged — and since it looked promising, he invited them to the call before releasing it; the forum cannot be used for advertising. He expects the post can now be released.

Discourse, the forum software, holds a new user's early posts for moderator review on several triggers, one of which is a post entered implausibly fast — a common sign of pasted spam. That matches what Stephen described. Inviting a vendor to present on the call before approving their post is a sensible way to separate a genuine contribution from advertising.

“Renting a certified UI” describes the common model in which an EHR embeds an e-prescribing vendor's own screens — often in an iframe — rather than building the workflow natively against an API. Weno's EZ Integration, for example, embeds Weno's DEA-compliant screens in OpenEMR via iframes. An API integration like LogiCoy's takes more work but gives a native workflow inside OpenEMR, which is what was being praised.

E-Prescribing Options at a Glance

Compiled after the meeting from the OpenEMR ePrescribe wiki page, earlier meeting summaries, and vendor websites.

Option	How it works with OpenEMR
NewCrop (now Ensora Health)	Long-standing integration provided through authorized OpenEMR vendors such as MI-Squared, which has maintained a DEA certificate for each OpenEMR iteration. Subscription-based; EPCS is an optional add-on requiring identity proofing. Routed over Surescripts. Relied-upon software on OpenEMR's 2017 certificate.
Weno Exchange	Bundled module with a no-code “EZ Integration” for OpenEMR 7.0.2 and later. Low cost: a sign-on fee to leave test mode, and most prescribers pay under \$200 a year. Uses Weno's own pharmacy network; pharmacies not on it receive a fax alert and an invitation to join. Coverage gaps reported in some areas (9/22).
DoseSpot	Not bundled. An embeddable, Surescripts-routed service with EPCS; Sherwin Gaddis built an integration module for a client earlier this year (9/22).
LogiCoy eRx	New. Module in development by LogiCoy, using its API over a direct Surescripts connection. ONC-certified standalone product (CHPL 10944) with PDMP integration, ID.me identity proofing, and RTPB; ePA in progress. A possible path to OpenEMR (b)(3) certification.

Release Timing and Review Volume

- **Next release:** expected next week, or possibly the week after, with as many pending fixes included as possible.
- **Billing fixes:** appreciation for the billing fixes that have been coming in.
- **Volume:** submissions have not yet plateaued, and reviewers are prioritizing small PRs to keep up.
- **Open PRs:** the aim is to keep the number of open PRs below 300. One person stepped away for two days and came back to find it at 300.

Three hundred open PRs is a lot for a project with a handful of active reviewers. The tactics described across recent meetings form a consistent picture: AI review (CodeRabbit) as a first pass, small PRs first, a standing PR slot in the weekly call, and the structure-first policies that make AI-assisted PRs easier to review. Size-first triage has a cost worth naming, though: large but valuable PRs — like the 5,000-line LBF module — can wait indefinitely unless someone deliberately picks them up, which is exactly what the new PR slot is for.

PR 13795 — LBF Statements Comes Into Core

Pull request: <https://github.com/openemr/openemr/pull/13795>

Where It Stood

- Simon recapped: the question was whether OpenEMR wants to bring in modules like this; opinion was split last time, and the decision was postponed.
- Brady's recollection: the issue was that this is a core feature. If it is a module, it would be a core module that is always on.
- Simon's one-line description: it takes data values and ranges and turns them into sentences in a text box.
- **A CodeRabbit aside:** the PR has accumulated many CodeRabbit reviews. Simon has had to resume CodeRabbit repeatedly and has been hitting its rate limit. The tip offered: CodeRabbit does not review draft PRs, so keeping a PR in draft while still iterating avoids using up reviews.

Background from the 9/15 summary: the module generates prose sentences from numeric values already recorded on a layout-based form — numeric bands, ratios between fields, and list or text tokens — and writes the paragraph into a textarea on the same form. It is about 5,000 lines across 33 files, with isolated unit tests. Earlier comments on the PR thread had questioned the maintenance burden and asked why Nation Notes or Text Notes could not do the job; Simon replied that neither reads other fields on the form, and added a test plan.

CodeRabbit's documentation fills in the details. Draft PRs are skipped by default, and marking a PR ready for review triggers the first review. Incremental reviews pause automatically after a set number of reviewed commits — five by default — and resume with “@coderabbitai resume” or a review request. Rate limits are per developer, per hour, and each push that triggers a review counts against them, so pushing several commits together rather than one at a time stretches the allowance. “@coderabbitai rate limit” reports what remains.

How the Module Works

- **Self-contained:** it uses its own tables and does not step on anything else. Built straight into the codebase, it would have been acceptable without any setting to turn it on or off.
- **Rules are set per LBF form,** in a dedicated menu for the module, not in the layout editor itself. An administrator builds the form as usual, then uses the module's interface to add rules to that specific form.
- **Not in the layout editor's “?” column:** Simon tried building the rules into the existing per-row conditions in the layout editor but did not get far, and concluded a standalone interface was better — though he is open to correction.
- **Output:** the generated sentences go into a single statements-and-conclusions text box at the bottom of the form, which can be customized — into separate text boxes, for example.
- **Across rows:** rules can span multiple rows, which the existing “?” conditions cannot. Brady described those conditions as fairly basic — hide this field, set a value — and crowded. Cross-row rules open the door to calculations: if these numbers add up past a threshold, put specific text at the bottom of the note.

- **Brady's view on the separate interface:** go with it for now, and revisit if people find themselves bouncing back and forth. Folding it into the crowded “?” column could lose the cross-row capability, and with AI it would be easy to migrate to a better, unified place later.
- **Jerry's angle:** his questionnaire work aims to make templates, questionnaires, and assessments assignable by context — encounters only, dashboard only, or both — and LBF could be part of that, since portal templates already render LBF forms. He wants to understand what differentiates Simon's work from what already exists, and said he had not yet had time to look closely.

The “?” column is the skip-conditions feature of OpenEMR's layout editor: for each field, an administrator can enter conditions referencing other fields, with an action, typically to hide the field or set a value. It is per-field logic about what appears on the form. Simon's module works at a different level: it reads the values of several fields after entry and writes a narrative from them. The two are complementary rather than overlapping, which supports keeping a separate interface.

There is a FHIR analogue relevant to Jerry's questionnaire work. FHIR Questionnaires support conditional display (enableWhen), and the Structured Data Capture guide adds computed fields (calculatedExpression) and extraction rules that turn answers into chart data. If questionnaires become a main form-building path in OpenEMR, cross-field calculations and narrative generation are capabilities the two efforts will eventually want to share.

Code Notes From the Live Review

- **Schema cleanup:** Brady reviewed the module's SQL on screen. Its table file creates the tables and then alters one — adding a match-token column and then modifying it — guarded by “if missing column” checks. Since the column is already correct in the table definition, those statements are unnecessary: fold everything into the CREATE TABLE so the schema is clean. Brady's guess, which Simon accepted, was that they were a development-time upgrade path from before the column existed.
- **Eric Stern's read:** the code is very well structured and the tests are good. Moving it into core would be straightforward: the source classes and templates would simply move, leaving a couple of public endpoints, which he would prefer to fit into the API structure rather than the interface directory.
- **Housekeeping:** a reminder note to be added to the PR.

OpenEMR's SQL upgrade files use special directives such as #IfMissingColumn and #IfNotTable so that an upgrade script can run safely against databases at different versions — each change is applied only if it has not been already. They are the right tool in upgrade files but redundant in a fresh-install definition. If the feature moves into core, as decided below, its table definitions would go into sql/database.sql for new installations and into the next version's upgrade file for existing ones. The preference for API endpoints over the interface directory reflects the front controller work discussed in recent weeks: new endpoints should be routed, not added as standalone PHP files under interface/.

Core or Module — The Policy Question

- **Brady's framing:** built into the codebase, this feature would be accepted, and nobody should be penalized for building it as a module. The project is trying to stop bundling third-party modules, but this is different — a module that is a core feature. The question is whether it should need turning on at all, or be always on like the other core modules.
- **Eric Stern:** if something is always on and always bundled, the module structure adds nothing, and it makes testing and installation trickier to follow. The test is applicability: if it applies to 90% of users, it belongs in core — still meeting all the testing standards.
- **Stephen Nielson:** it depends. Some installations want to trim feature sets to improve their maintenance and security posture. The key distinction is between foundation, infrastructure, and core components on one side, and feature sets people would turn on and off on the other. Something that is on everywhere, all the time, only gains abstraction overhead from being a module; something people might reasonably turn off should be modular, even if it is usually on. Either way, it should stay cohesive and isolated so it is testable.
- **The phrase that summed it up:** if it is always on and cannot be turned off, it is not a module — it is “core in module clothing.”
- Stephen Nielson then had to leave.

- **Afterward:** Brady was unsure where Stephen had landed, since he has favored modules in the past. Eric: anything vendor-specific — the e-prescribing module from earlier in the call, for example — should be its own Composer package, even if included by default. It was also pointed out that if this feature is core, its separate database file should not exist; the table definitions go with the normal migrations into database.sql. Eric: with no reason to turn it off, put it directly in core and skip the module packaging complexity, which fits better with testing — database testing with modules is awkward. Where a feature is off by default, or there is good reason to turn it off — some installations do not need e-prescribing — moving toward a module makes sense.
- **Where Brady and Eric converged:** this is a value-add for layout-based forms with no reason to turn it off. It should live next to LBF, which is in core today; if LBF eventually moves into a module, this moves with it. It could also offer a path for migrating the existing LBF code — described as spaghetti at this point — since it arrives with a clean structure, provided it uses the modern architecture.

The Decision

- **Bring it straight into core,** not as a module — in something like src/Services/LBF — with the goal of eventually restructuring the existing LBF code the same way. Simon was asked whether he was comfortable with this, so the question need not be revisited.
- **The separate rules interface stays for now.**
- **Jerry was fine with it,** and will look at how it relates to his questionnaire and template work.

This resolves a thread running since 9/15, when the PR was deferred amid disagreement about modules versus core. On 9/22 it was awaiting a module-policy follow-up; on 9/29 Eric gave his order of preference for extending OpenEMR. Today's discussion produced a practical three-way rule: if a feature applies to nearly everyone and there is no reason to turn it off, it goes in core; if it is optional — off by default, or something installations may reasonably disable — it is a module; and if it is vendor-specific, it is a separate Composer package, even when bundled. That is consistent with the 9/15 clarification that applicability, not packaging, decides what belongs in core.

Brady's remark that the code could provide a path for the legacy LBF code echoes the 9/15 argument for the module form — that an isolated, well-structured piece could become the model for eventually extracting LBF. The decision keeps that benefit by placing the code in src/ as a service, where modern structure and tests can sit alongside the legacy LBF code and gradually absorb it.

CCDA, Node, and What Jerry Does Next

- **A shout-out:** Jerry got a huge shout-out on the call for getting the CCDA work over the line — huge progress over the past couple of weeks. Being able to turn off the Node infrastructure will make the project much easier to work with, and more secure. Jerry was happy it worked out.
- **Next for Jerry:** either the audit work or the questionnaire manager. The questionnaire manager has to be in 8.5: today, any FHIR questionnaire imported into the repository appears on the dashboard, and the only way to manage it is to delete it from the repository. The questionnaires and the audit are his next mission.

On 10/6 the PHP CCDA implementation was close to merging, with a separate PR to remove the Node path to follow, and the CQM calculator — the last Node service — was Jerry's next target. The remark about turning off the Node infrastructure suggests the CCDA side is now complete; whether CQM has also moved was not stated. Once all three Node services (Schematron, CCDA, and CQM) are gone, Node remains only as a build-time tool for front-end packages, removing a separate runtime, its patching, and a class of supply-chain exposure from every installation. The questionnaire manager is the piece Jerry described on 10/6 as on hold behind the audit work; the 8.5 release now sets a deadline for it.

FHIR Writes — An Alert Instead of a Global

- **Brady asked whether the global gate for FHIR writes had been added.** Jerry implemented it as an alert instead: a global would have meant touching code all over the place, and it got messy.
- **Brady accepted that the alert covers it,** because automated (system) clients cannot get write scopes at all. Writes are possible only for clients an administrator must enable, and enabling one now brings up a prominent alert saying that writes will be allowed and listing which write scopes the client uses. Jerry kept the wording brief.

This revises last week's agreement. On 10/6 the group agreed on a global that would keep FHIR writes off by default, with a warning on enabling write-scoped clients as an optional extra, and Brady argued that the hard gate was what mattered. Today's reasoning is that the chokepoint already exists: writes are user-context only and require an administrator to approve and enable the client, so the moment of enabling is the natural place for an explicit, informed decision. The trade-off is that an alert depends on the administrator reading it, whereas an off-by-default setting protects installations whose administrators never consider the question. Jerry's own 10/6 point — that overloaded administrators may approve clients without researching them — applies to alerts too, which is one reason the audit-logging work remains important.

Settings Screen — Handling Unreadable Encrypted Values

- **Agreement reached:** after some back-and-forth on how the settings screen should behave when encrypted values cannot be read, the resolution was a loud banner. Both sides were right: an installation that reaches that state is in serious trouble, but there is no harm in handling the case.
- **The behavior:** when stored encrypted values cannot be read, the screen shows a prominent alert telling the administrator to stop and investigate rather than save settings. There is nuance: a single unreadable leftover value can be accepted and lost, whereas everything being unreadable points to a configuration problem.
- **Why it matters:** making the screen more reliable is a big win. At one point, a single item failing to load correctly led to every setting in the system being wiped.

A settings page typically submits every field when saved. If a value cannot be read and displays as blank, an administrator who saves the page — perhaps to change something unrelated — overwrites the stored value with the blank. Failing loudly, and discouraging a save, keeps a display problem from turning into data loss. The distinction between one bad value and all bad values is the right diagnostic: the first is residue, the second means something fundamental in the configuration has changed.

Luis Santos — Moving a Code Import Script to the Console

- About a month ago, Luis was asked to break up a code import script and move it into the console, along with another command-line area. He has tried, but invoking the new command fails with an “unable to resolve” error related to factory registration. He shared his screen to walk through the commits, which he modelled on how other commands are defined.
- **Suggested next step:** run the command with `-vvv` or `--debug` to get a stack trace, which should point directly at the registration problem. The reviewer offered to look after the call and leave comments on the PR rather than debug live.

OpenEMR has been consolidating its assorted command-line scripts into a single Symfony Console application, so that each command has consistent argument handling, built-in help, and a testable class. In Symfony Console, `-v`, `-vv`, and `-vvv` raise verbosity, and `-vvv` an exception prints a full stack trace. An “unable to resolve” error at that stage usually means the command, or one of its dependencies, has not been registered with the container or factory that constructs commands — a wiring problem rather than a logic one. The 10/6 meeting noted a multisite bug apparently introduced by the same command work, a reminder that moving scripts into a shared framework also changes their assumptions.

SMART and API Test Tools

- As Brady was leaving, Jerry reported that his local tools for the SMART app and the API were all failing, and asked whether they were working for Brady. Brady said the tests pass in CI: the SMART test walks through the whole flow, including writing a questionnaire and confirming the write, while the API test covers only initial authentication. Jerry thought he might have broken something in the codebase the day before, and the two agreed to follow up directly.

Automated SMART and API tests in CI are the kind of tooling Brady described on 10/6 — build the tests and let AI use them. A SMART test that writes a questionnaire exercises exactly the FHIR write path discussed above, so it doubles as a regression test for that work.

Items Discussed and Proposed

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

Area	Discussed / Proposed
LogiCoy e-prescribing	Introduction and demo of the standalone product and the OpenEMR module (basic prescribing done; ePA and other advanced APIs to follow). To be discussed by project leadership, with a follow-up meeting including LogiCoy's CEO and US sales lead.
(b)(3) certification	Desired but not required to proceed. Surescripts approval via a questionnaire sponsored by LogiCoy. Cost and usability-testing work to be assessed. Target is SCRIPT 2023011 and the new (b)(4) RTPB criterion.
Partnership model	LogiCoy to propose at the next call. A technology-partner category like EMR Direct suggested, rather than a support-vendor tier. Onboarding and billing design open.
LogiCoy forum post	Held by the forum filter; expected to be released.
PR slot	PRs to be reviewed around the middle of each call, aiming for one or two a week.
PR 13795	Bring into core directly, near LBF (e.g. src/Services/LBF). Fold the SQL into the table definitions and database.sql. Keep the separate rules interface for now. Long-term goal: restructure existing LBF code the same way.
Module policy	Always on with no reason to disable: core. Optional: module. Vendor-specific: separate Composer package, even if bundled.
Release	Expected in the next week or two.
CCDA and Node	Credit to Jerry. Next: audit work or the questionnaire manager, which must be in 8.5.
FHIR writes	Implemented as an alert when enabling a write-scoped client, rather than a global. Accepted.
Encrypted settings	Loud banner when stored values cannot be read; agreement reached.
Console migration	Luis to debug with -vvv; reviewer to comment on the PR.
SMART and API tests	Passing in CI; Jerry's local tools failing. Follow-up offline.

Notes

- **The connecting thread:** the two halves of the meeting turned out to be linked. The module rule settled in the second half — core for what everyone uses, modules for what is optional, separate packages for what is vendor-specific — answers a question the first half raised. An e-prescribing integration is vendor-specific and not every installation needs it, so by that rule it would ship as its own package even if bundled for certification. Packaging is a technical question; whether to bundle is a partnership and certification question.
- **E-prescribing has been a recurring gap.** On 9/22 a migrated client found Weno's pharmacy coverage inadequate, and DoseSpot was suggested as an alternative. LogiCoy's direct Surescripts connection addresses coverage, and its certification offers something neither existing integration currently gives OpenEMR: a path back to (b)(3) certification. The regulatory deadlines in 2027 and 2028 make the timing of any certification effort important.
- **Carried forward:** the items deferred for Stephen Nielson's return — the multisite, subdirectory, and non-Docker deprecation plan with Eric's migration guide, and procedure writes over FHIR — were not reached this week.
- Practitioners who e-prescribe in OpenEMR today — especially those prescribing controlled substances, or working around pharmacy coverage gaps — are best placed to say what a bundled e-prescribing option needs to do well. Anyone who has been through Surescripts onboarding or (b)(3) certification could also help the project judge the effort involved. Comments on the forum or at the weekly call are welcome.