



Why Prior Authorization Technology is Harder Than It Looks

A Provider's Guide to the Full
"Continuum of PA Complexity"



Why Prior Authorization Technology Is Harder Than It Looks: A Provider's Guide to the Full "Continuum of PA Complexity"

Written By: Michael Kolb, Chief Growth Officer, Humata Health

In 2019, many providers were asking a deceptively simple question:
Should we buy or build prior authorization technology?

The 'build' camp was particularly attractive at the time - health systems saw an opportunity to leverage their internal IT teams to create tailored solutions that could integrate directly into workflows and be maintained in-house.

Fast forward to today, and that enthusiasm has tapered off. Why? **Because prior authorization is one of the most multi-faceted, interconnected, and variable processes in all of healthcare.** It touches dozens of systems, depends on clean and timely data from multiple sources, and requires deep knowledge of payer rules, coding logic, and clinical documentation. Building it internally isn't just an IT project - it's a full-scale operational transformation that demands expertise across revenue cycle, clinical operations, integration engineering, and AI-driven automation.

This guide offers a transparent look at the full continuum of considerations for anyone thinking about building or buying a prior authorization solution.

CONTINUE



Integration or Bust

Providers increasingly prefer solutions that integrate natively into their EHR - especially Epic. The expectation is that solutions live inside the EHR interface, working seamlessly alongside core patient access activities such as eligibility, estimates, and registration.

But “integration” isn’t just about a window inside the EHR. To truly solve prior authorization, technology must have deep, bi-directional connectivity into the full continuum of internal and external data sources:

- Internal Sources
 - Clinical Orders & Referrals
 - Scheduling & Encounter Data / Patient Appointment Records
 - Clinical Documentation (Internal/External)
 - Document Management Systems
 - Coding & Charge Capture / Diagnosis & Procedure Coding Data
 - Claims & Remittance Feeds / EDI 835/837 Transactions
- External Sources
 - Payer Coverage & Medical Necessity Guidelines
 - Health Plan Portals / TPA Authorization Portals
 - Payer Integration APIs / Plan Connectivity APIs
 - Third-Party Vendor APIs / External Partner Integrations

This level of integration is non-negotiable. Without it, prior authorization remains fragmented, staff continue to swivel-chair across systems, and errors creep in - leading to delays, denials, and poor patient experience.

It’s also critical to recognize that integration is not a one-time build.

Payer portals change constantly, medical necessity guidelines are updated regularly, and new federal mandates (like WISeR in 2026, CMS 0057 in 2027, etc.) will reshape workflows. A sustainable solution requires ongoing maintenance, optimization, and extensibility to keep pace with these changes.

Key Needs: EHR integration knowledge and capabilities; bi-directional connectivity across orders, appointments, clinicals, coding, policies, payers, and delegated entities; commitment to continuous maintenance



The Hidden Complexity of Orders & Auth Requirements

Triggering prior authorization is not as simple as “order placed, auth begins.” In reality, the process is riddled with variation:

- Some health systems generate referrals automatically when an order is placed
- Others create “shell referrals” with no order at all, introducing gaps
- Many orders lack CPT codes until after front-end coding is complete — and without service-level data, automation can’t reliably start
- In Epic and other systems, different modalities (referrals, auth-certs, series encounters) create further divergence

Mapping the Maze

Technology must adapt to these realities. That means mapping orderables to CPT codes, linking orders to appointments for site-of-service considerations post-scheduling, and integrating seamlessly with both payers and delegated entities to confirm requirements. To scale, this requires connectivity not only to national payers but also state/regional carriers and TPAs.

Payer Connectivity at Scale

A payer connectivity strategy is more than APIs. While some payers and delegated entities offer them, many still rely on portals. That means technology must:

- Navigate portals with RECAPTCHAs and multi-factor authentication
- Capture and store screenshots of determinations, which are vital for appeals
- Orchestrate a multi-modal approach that blends API calls, portal technology, and human-in-the-loop escalation when needed

Submission Entity Routing

One of the biggest operational headaches for providers is determining where to submit. Payers frequently delegate to TPAs, creating confusion over which entity owns the requirement check and the submission. A “waterfall” logic is essential: query the payer first, then seamlessly redirect to the delegated entity when appropriate. A sustainable solution must be designed for continuous evolution: ongoing iteration, monitoring, and optimization are non-negotiable to stay compliant and efficient.

Key Needs: Deep EHR integration, payer connectivity via API/portal, CPT mapping, delegated entity routing, multi-modal requirement retrieval

Cracking Medical Necessity

The lifeblood of authorization approval is the payer's medical necessity policy - but these aren't centralized. They're scattered across websites, portals, and specialty-specific guidelines.

Building a robust solution requires:

1. API access to retrieve policies
2. Integration into the EHR so they're visible at the point of authorization
3. Generative AI to break them into actionable, structured criteria sets

With this foundation, technology can mine clinicals in real-time to ensure criteria are met before submission. This isn't just about getting an authorization - it's about getting paid. Just because you receive a prior authorization does not guarantee reimbursement. If the clinical documentation submitted doesn't align tightly to the policy's requirements, the claim can still be denied after services are rendered.

A strong build should therefore connect the policy requirements directly to the documentation process, ensuring every relevant detail is addressed and supported before the patient even arrives for their procedure. This mitigates both front-end and back-end denials, protecting revenue and reducing rework.

Key Needs: EHR integration, payer policy APIs, generative AI for policy decomposition.



just because
it's authorized
doesn't mean
it will be paid

The AI Bridge: Connecting Documentation to Authorization

One of the most labor-intensive tasks in prior authorization is clinical bundling — identifying which documents justify the request. These live in multiple locations: chart review, media tab, Epic CareEverywhere, and external document management systems like OnBase.

A modern solution must:

1. Pull documents via API or virtual machine (VM) integrations
2. Use NLP, OCR, and document classification to become familiar with the clinical content
3. Leverage a dynamic learning model to recommend the right documents based on plan, CPT, and diagnosis
4. Continuously adapt from staff interactions to alter recommendations over time

Including the right documentation the first time has a massive impact on reducing peer-to-peer requests - one of the most dreaded disruptions for physicians. **By using technology to intelligently guide which documents are sent to the payer, providers can reduce unnecessary escalations, protect physician time, and improve the likelihood of a first-pass approval.** This not only streamlines the process but also strengthens provider-payer relationships by improving submission quality.

Key Needs: API integration, VM integration, NLP, OCR, dynamic learning model



Integrating with and interpreting clinical documentation is critical



Scalability across specialties, payers, and process is the key to maximize value

Submission: Solving for Today While Preparing for Tomorrow

The holy grail of prior authorization is truly touchless submissions - but the industry isn't there yet. While entities like eviCore and Availity have opened APIs that enable electronic submission, scale is limited. Da Vinci-compliant APIs tied to CMS 0057 may unlock broader adoption by 2027, but in today's environment, most payers still rely on portals - and some even require fax.

A robust technology solution should leverage the automated clinical documentation to then:

1. Auto-submit where it's feasible based on today's connectivity
2. Automate coversheet population and eFax submission
3. Augment portal workflows to reduce manual burden
4. Auto-answer attestation questions using generative AI, mapping them to the clinical indications in the order

In addition, a truly comprehensive build must account for secondary and tertiary coverage. Many patients have more than one plan, and failing to secure authorizations for all applicable coverages can lead to payment delays or denials.

Lastly, one often-overlooked opportunity in prior authorization routing is leveraging the alpha prefix from a member's insurance ID - particularly for large payer families like the Blues. Today, many organizations default to a catch-all category like "Blues Other" for out-of-region members, because the routing logic isn't mapped down to the three-letter prefix. This means high-volume traffic from states like Texas or Illinois is handled manually or sent through generic workflows, adding friction and delay. If you can decode those prefixes and automatically direct each request to the correct portal, you not only remove a major bottleneck but also unlock significant operational efficiency.

Key Needs: Multi-modal submission (API, portal, eFax), generative AI for attestations, portal workflow automation, secondary & tertiary coverage handling, alpha prefix routing logic, scalability for future standards (Da Vinci, CMS 0057)



Code-Level Clarity: Statusing Every Authorization

Today, most staff still log into payer portals manually to check status updates - a time-consuming process that creates unnecessary touches.

A modern solution must automate statusing across all high-volume payers, using a multi-modal approach (API where available, portal automation where APIs don't exist, and even outbound call automation for high-value or urgent requests).

Crucially, statusing must occur at the CPT-code level, not just at the order level. A single prior authorization request often covers multiple CPT codes, and payers may return different statuses — approvals, partial approvals, denials — tied to each code. If the system only tracks at the order level, critical updates will be missed.

Technology must also normalize and integrate the full authorization data back into the EHR, including: authorization numbers, from/through dates, number of visits approved, and any denial or "additional info needed" messages.

Key Needs: multi-modal statusing (API, portal, outbound calls), CPT-level multi-auth tracking, integration of auth numbers/dates/visits into EHR, digitization of payer messages, scalability across all payers



Preventing Denials Through Real-Time Data Monitoring

In prior auth (and in healthcare in-general), data changes constantly - eligibility updates, scheduling shifts, clinical indications. Without constant monitoring, these changes can “break the auth”.

Technology should integrate logic to flag updates to CPT, coverage, diagnosis, or procedure dates - and tying them to referral statuses - providers can prevent retro denials and reschedule delays.

In addition, experimental codes, incorrect diagnoses, and payer-specific quirks can all be caught earlier with mapping exercises and workflow design that flag risk scenarios before submission.

Key Needs: EHR integration, reconciliation model, integration of 835/837 data



CPT Code Accuracy: Closing the Gap

A common cause of denials: what’s ordered is not what’s rendered.


- Predictive models can flag “likely CPT changes” at the point of auth
- Post-service integration with computer-assisted coding (CAC) can identify final codes earlier (prior to abstraction) to route for retro
- Data from these mismatches can be fed back to physicians for ordering education

Key Needs: EHR integration, ML model, coding integration

Making Ambient Dictation Work for Prior Auth

Ambient clinical documentation technology is booming - and there's no shortage of industry buzz about enabling prior authorization at the point of care. The concept is compelling: as the provider speaks with the patient, documentation is generated automatically, and prior auth criteria can be addressed in real time.

The reality, however, is more complex. While ambient tools can capture rich, detailed encounter notes, most prior authorizations require far more than what comes from a single visit or specialty. They demand a bundle of records that span multiple specialties, encounters, and sometimes even external facilities. In other words: a transcript alone isn't enough. True success requires deep access to the broader clinical record.



To operationalize prior auth at the point of care, technology must:

- Time CPT code availability accurately to align with guideline requirements
- Decompose medical necessity guidelines with generative AI into the exact criteria payers require
- Aggregate all relevant documentation across EHR modules and external repositories
- Leverage NLP, OCR, and classification models to surface and package documents automatically
- Trigger authorizations intelligently, accounting for payer-specific requirements like scheduling location

Without this broader integration, point-of-care prior auth risks being an isolated gain instead of a true end-to-end automation. The future lies in pairing real-time documentation with full clinical data aggregation — enabling providers to close the loop in the moment, with confidence that every payer requirement is met.

Key Needs: medical necessity decomposition, clinical documentation integration, NLP/OCR, multi-specialty document aggregation



Positively Impacting Scheduling Through Interconnected Authorization Technology

With the right analytics, providers can schedule smarter - using real turnaround times for specific plans and procedures to reduce reschedules and improve utilization.

Today, most providers schedule a fixed number of days out by specialty, based on historical practice rather than current payer performance. This can create gaps in two directions:

- Scheduling too far out may delay care unnecessarily, frustrate patients, and reduce throughput
- Scheduling too soon risks not having the authorization in time, which leads to reschedules, lost revenue, and negative patient experience

When using technology for prior auth, metrics should track “days to decision” and “days to approval” at both the plan and CPT code/code set level. By combining these metrics with scheduling logic, providers can:

- Precisely match appointment lead times to actual payer behavior
- Reduce the volume of rescheduled appointments
- Maximize schedule utilization and revenue capture
- Improve patient satisfaction by avoiding last-minute changes

Scheduling optimization should be an operational output of prior auth analytics, not an afterthought - ensuring that the entire process, from order to appointment, is aligned with real-world timelines.

Key Needs: Analytics at plan and CPT/code set level; integration of metrics into scheduling workflows

Analytics & Ongoing Maintenance

Going live is not the finish line - it's the starting line. Without deep KPIs such as peer-to-peer rates, first-pass approvals, days to decision, denial root causes, and reschedule rates, providers can't monitor performance or identify automation opportunities.

But analytics alone aren't enough - ongoing maintenance is essential for sustained success. Prior authorization is a constantly shifting environment:

- Authorization requirements change regularly, sometimes with little notice
- Payer portals are frequently updated, which can break existing automations or require reconfiguration
- Volumes fluctuate, which may require dynamic load-balancing of portal credentials to ensure performance
- Medical necessity guidelines are updated on an ongoing basis, and these must be integrated and decomposed quickly to avoid compliance gaps



This means any solution must have a dedicated, full-time maintenance and optimization process. Teams should continuously review metrics, denials, and exception cases to identify both quick fixes and longer-term enhancements. Without this discipline, even the most sophisticated build will degrade over time - leading to increased manual work, more denials, and reduced ROI.

A high-performing PA technology stack is as much about ongoing tuning and optimization as it is about the initial go-live.

Key Needs: Analytics via Tableau, Power BI, or similar; dedicated maintenance resources; proactive change management processes



The Bottom Line

Building prior authorization technology is not a coding exercise — it's a complex orchestration of clinical, operational, technical, and payer-specific expertise. Every integration point, every data element, and every decision impacts whether an authorization is approved on the first pass or delayed for days.

The 2019 “buy vs. build” conversation made sense in a simpler era. Today, the question is less about whether you can build it and more about whether you have the multidisciplinary depth to build, maintain, and continuously adapt it in a shifting payer and regulatory landscape.

This isn't a feature that can be bolted onto broader EHR platforms. **Prior authorization demands complete focus and dedicated innovation from those committed to solving it end-to-end. Without that focus, the complexity and variability of payer rules, clinical documentation, and submission workflows will overwhelm even the best-intentioned solutions.**

Because in prior auth, complexity isn't the exception - it's the rule.

Ready to see how Humata can solve your PA problem?

Contact us today at info@humatahealth.com for an interactive demo.