Prior Authorization Use Case
AI Agent with Guardrails

GRANITE FORT
ADVISORY

AI Transformation, Governance, Risk & Compliance
Clarity. Compliance. Confidence.

EXECUTIVE SUMMARY

Artificial intelligence (AI) is moving beyond predictive analytics into operational transformation. In healthcare, where margins are thin and regulatory scrutiny is high, the rise of AI agents, autonomous, persistent systems that act within defined boundaries, marks a step change.

These agents are not static chatbots or dashboards. They can read, reason, and act while operating under policy and compliance guardrails that ensure safety and trust.

This paper explores how AI agents are reshaping healthcare, with a spotlight on one of the industry's most pressing administrative challenges: **Prior Authorization (PA)**.

We show how an **AI Prior Authorization Agent with Guardrails** can reduce administrative burden, speed access to care, and increase compliance while maintaining transparency and accountability.

AI AGENTS
IN HEALTHCARE

Providers, payers, and patients are overwhelmed by the volume of data, paperwork, and rules that govern every aspect of care. AI is beginning to address these challenges, moving from back-office analytics to front-line automation. What differentiates the current wave is the emergence of AI agents.

WHAT ARE AI AGENTS?

AI agents are autonomous systems designed to sense, reason, and act on information in real time. Unlike task-specific bots or rules engines, they combine multiple capabilities:

- Continuous monitoring of structured and unstructured data sources.
- Document reading and comprehension of medical records, lab results, and clinical notes.
- Decision-making within guardrails defined by policies, regulations, and organizational standards.
- Task execution such as submitting forms, scheduling, or sending alerts.
- Escalation to humans when ambiguity, exceptions, or risks arise.

USE CASE:
PRIOR AUTHORIZATION AGENT

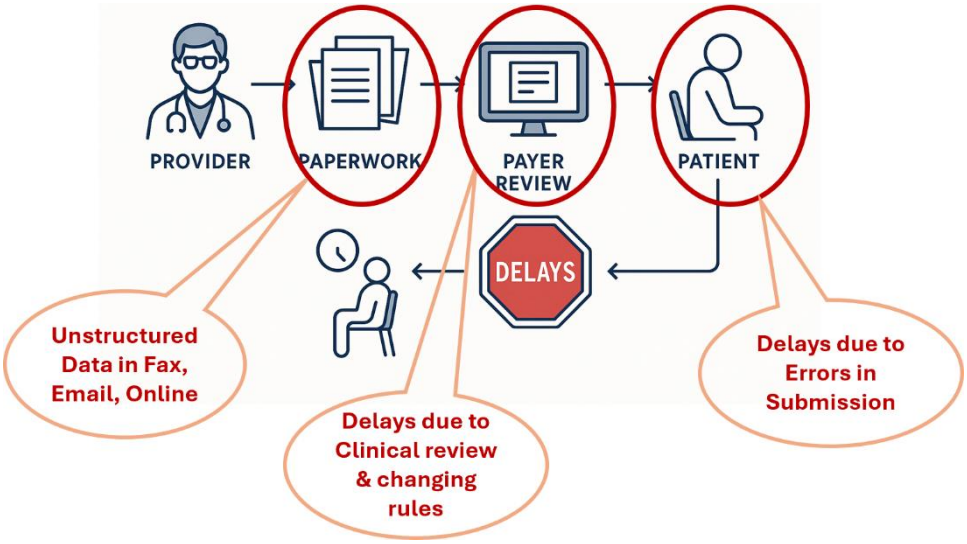

Prior authorization (PA) is one of healthcare's most criticized processes. Designed to ensure medical necessity and control costs, it has become a barrier to timely care and a source of tension between providers, payers, and patients. The result is a system that is inefficient, costly, and often unsafe.

An AI Prior Authorization Agent with Guardrails reimagines the process by automating the intake, validation, submission, and tracking of prior authorization requests, but doing so within a framework of embedded safeguards that preserve compliance, fairness, and human oversight.

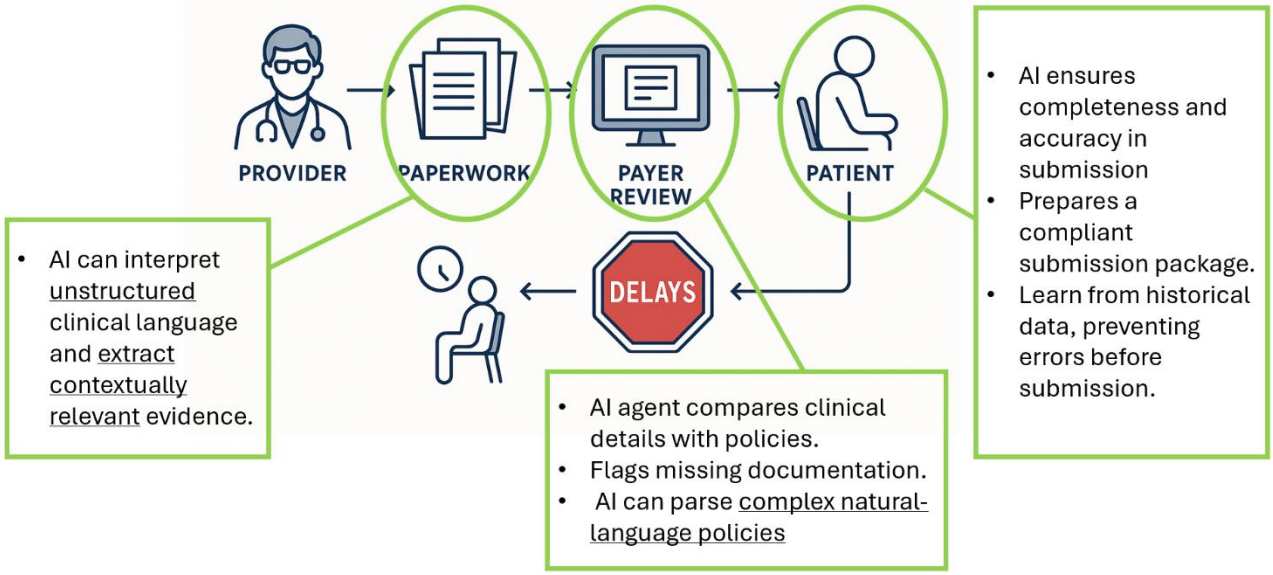

A. Reads and Understands Documents

The agent ingests medical records, lab results, imaging reports, referral letters, and scanned faxes, converting them into structured data using NLP and OCR.

Unlike traditional automation, AI can interpret unstructured clinical language such as 'failed metformin' or 'HTN' and extract contextually relevant evidence.

Data handling complies with HIPAA standards, ensuring that sensitive patient information is protected at every step.

B. Applies Policy Knowledge

The agent compares extracted clinical details with payer-specific medical necessity policies. It flags missing documentation before submission.

Unlike static rules engines, AI can parse complex natural-language policies and adapt dynamically to changes.

The agent cannot override or reinterpret payer medical policies. It only applies published criteria, ensuring clinical integrity.

C. Validates and Prepares Submissions

The agent ensures documentation is complete, coding is correct, and formatting matches payer specifications. It prepares a compliant submission package.

AI models learn from historical denial data, predicting errors and preventing them before submission.

Validation logic is locked to payer requirements. Ambiguous cases are routed to human review, ensuring accuracy and fairness.

D. Submits and Tracks in Real Time

The agent submits requests directly to payer portals or APIs and monitors their progress, alerting staff when responses are delayed or deadlines are near.

Instead of static tracking, AI enables adaptive workflows, prioritizing urgent cases and generating plain-language status updates.

Every action is logged, creating a transparent audit trail for payers, providers, and regulators.

E. Escalates with Guardrails

The agent never operates beyond its boundaries. Complex, ambiguous, or high-risk requests are automatically escalated.

Policy: Enforced adherence to medical necessity criteria.

Regulatory: Compliance with CMS and state-mandated timelines.

Human-in-the-Loop: Clinicians remain decision-makers in cases requiring judgment.

Auditability: Every decision is explainable and retrievable for appeals or compliance reviews.

STRATEGIC IMPLICATIONS: **WHAT MAKES THIS DIFFERENT**

This is more than automation. It is agentic AI, governed by compliance guardrails. Unlike traditional rules-based systems, it can interpret unstructured clinical language and adapt to evolving policies.

Unlike unchecked AI, it cannot act outside payer rules, regulatory mandates, or human oversight. The result is not just efficiency - it is a system that strengthens trust across payers, providers, regulators, and patients.

As the industry moves into an era of agentic AI, the central question will shift from 'Can we automate this?' to 'How do we ensure automation is safe, compliant, and trusted?'

The answer lies in agents with guardrails, scalable, explainable, and aligned with the mission of healthcare.

LOOKING AHEAD
THE PATH TO AI SUCCESS

Successful AI requires strategy, governance, and continuous transformation.

Need expert guidance to align your AI initiatives with measurable business impact?

Contact us to schedule your AI Strategy Review today.

Granite Fort Advisory
Dallas, TX, United States
Tel: +1-469-713-1511
Engage@GraniteFort.com
www.granitefort.com

GRANITE FORT ADVISORY

AI Transformation, Governance, Risk & Compliance

Clarity. Compliance. Confidence.