OSP

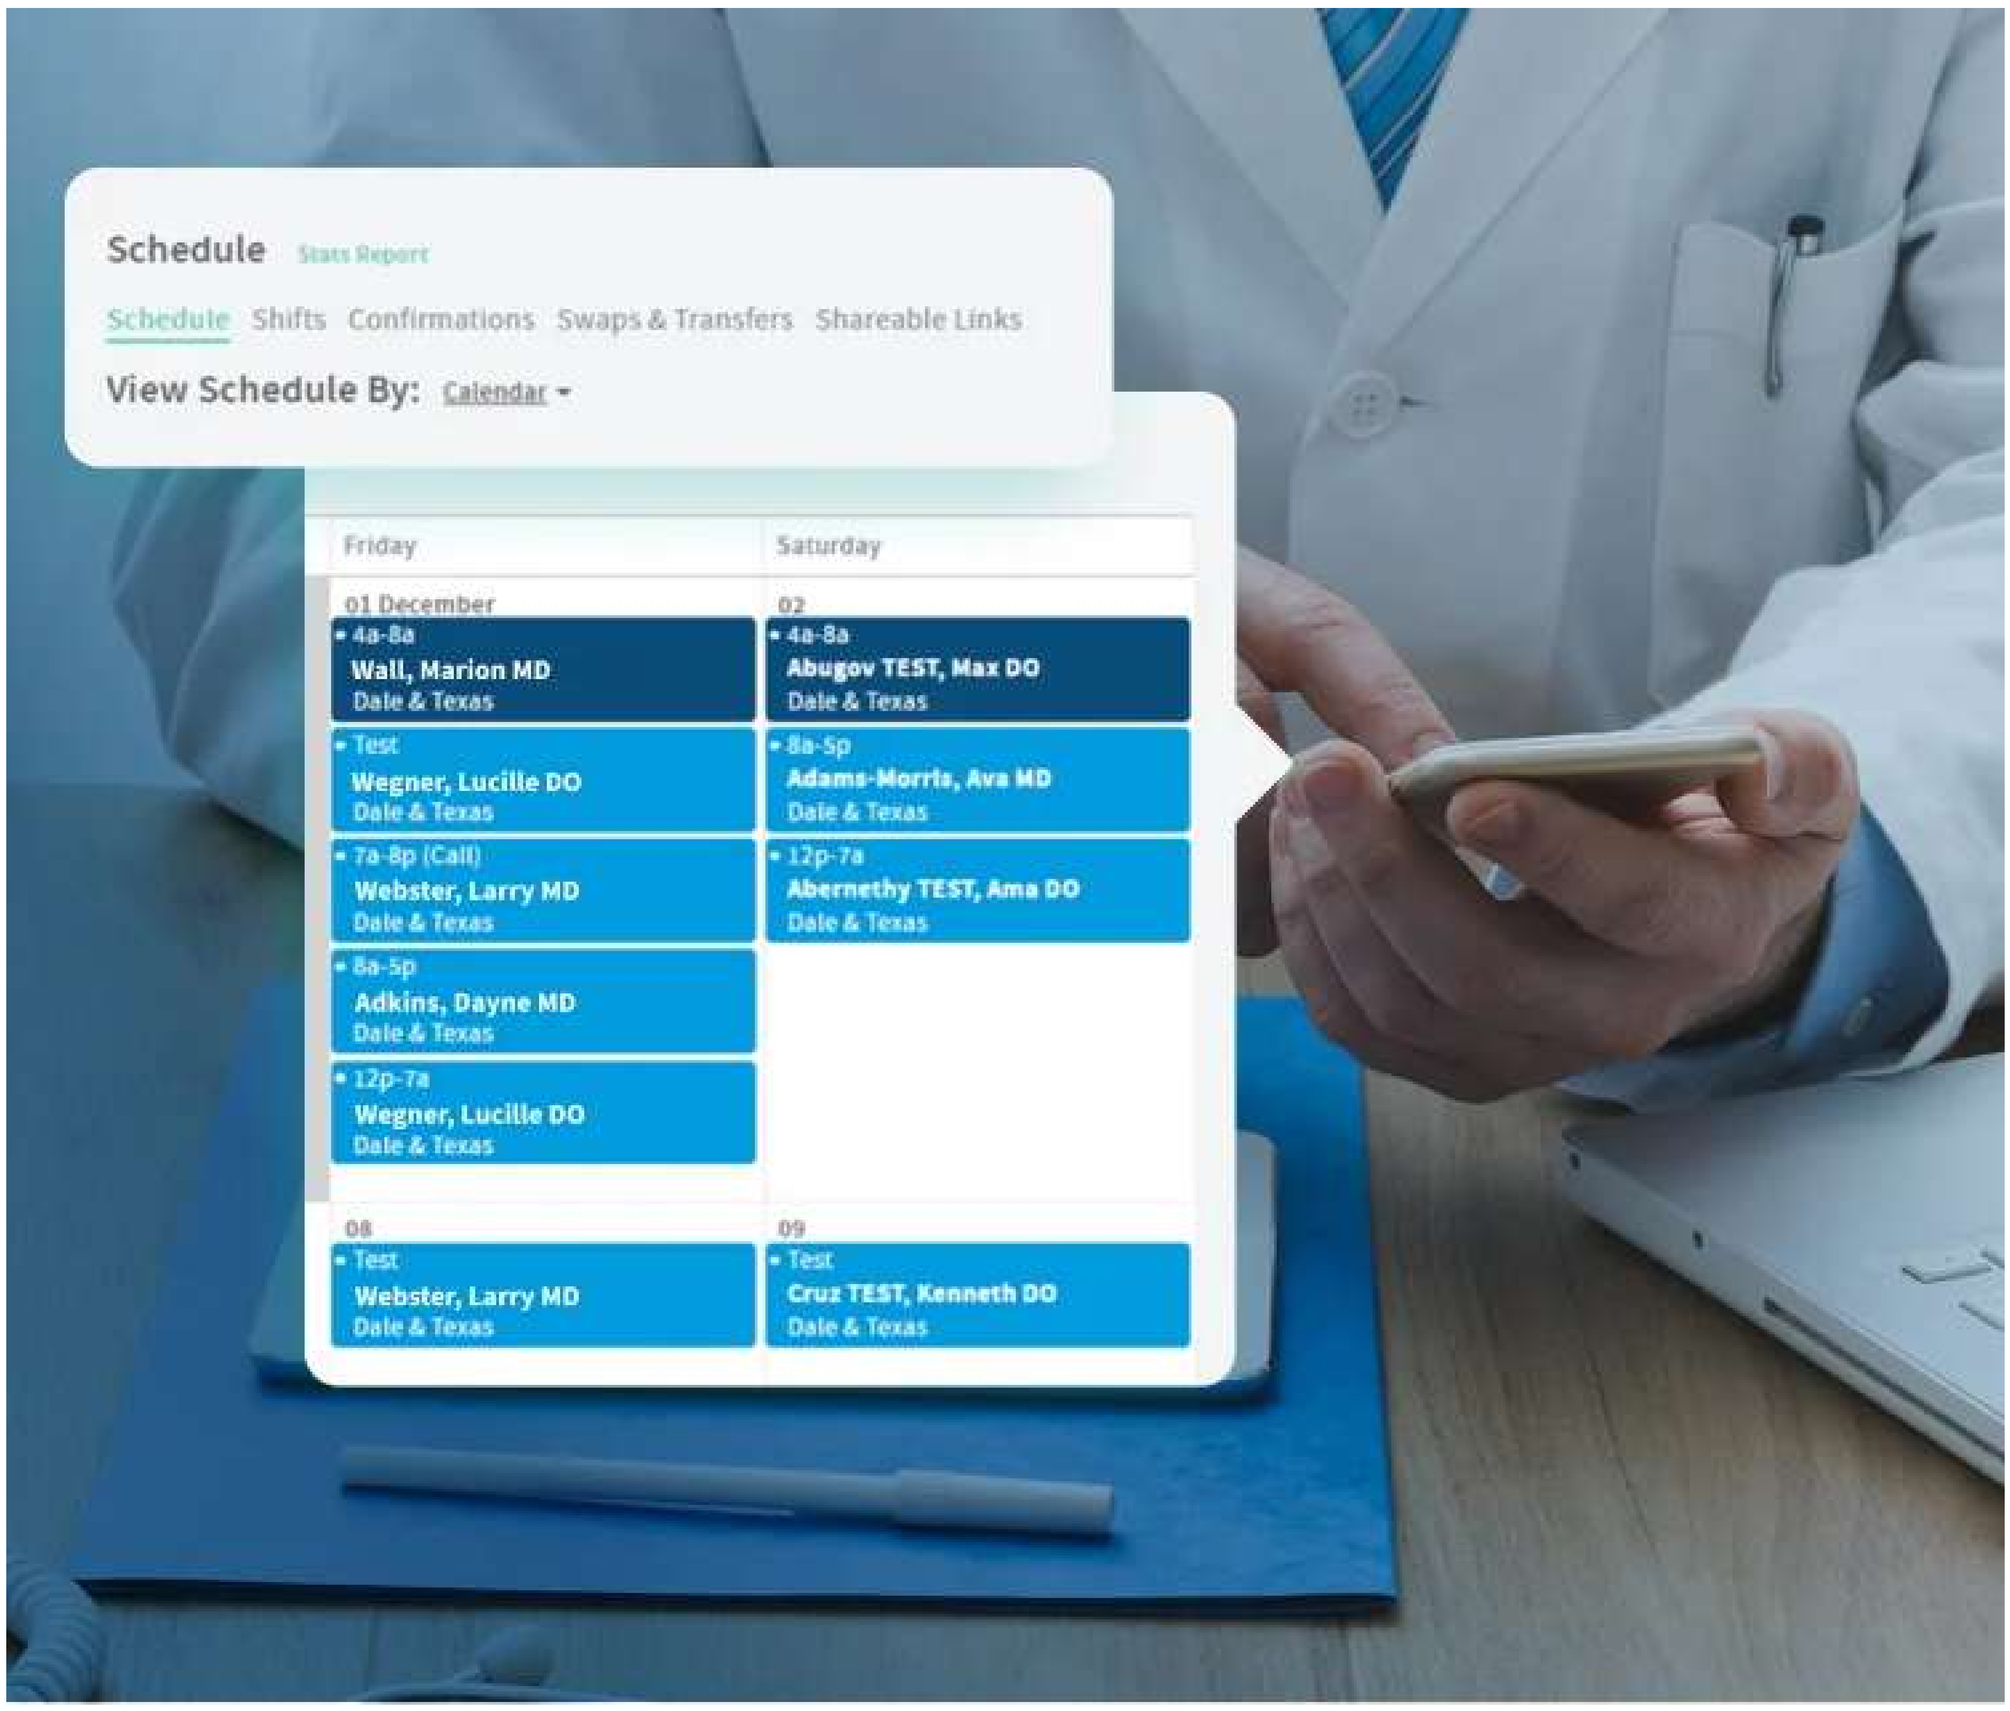

INTELLIGENT CLINICIAN SCHEDULING

AI-Augmented Optimization Using
Linear Programming

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Hospitals today face major challenges in scheduling the right number of clinicians at the right time. While predictive tools can forecast patient demand, most healthcare systems still rely on manual or outdated methods to create staff schedules. This leads to overworked clinicians, last-minute shift changes, and gaps in patient care.

This white paper presents a smarter approach: combining **Linear Programming (LP)** with **AI agents** to build optimized, flexible staffing schedules. LP helps solve complex problems involving multiple constraints— like clinician availability, work-hour limits, and personal preferences—while AI agents add flexibility, automation, and usability.

We outline how this solution works, how it connects with systems like Epic, Workday, or UKG, and how it can be tailored for real-world hospital environments. Our goal is to help healthcare leaders like you bridge the gap between demand forecasting and workforce planning—accurately, efficiently, and at scale.

THE HEALTHCARE STAFFING CRISIS

The clinician shortage in the U.S. is reaching critical levels. According to the Association of American Medical Colleges (AAMC), the country may face a shortfall of over 1**24,000 physicians by 2034,** with nursing shortages compounding the issue in nearly every state. At the same time, **burnout rates are surging,** with over **60% of U.S. physicians** reporting emotional exhaustion, primarily driven by unmanageable workloads and erratic scheduling.

These workforce challenges are amplified by:

- **Rigid scheduling systems** that don’t adapt to changing patient volumes
- **Lack of integration** between predictive analytics and workforce planning tools
- **Inability to account for clinician preferences,** leading to dissatisfaction and high attrition
- **Operational silos** between HR, nursing leadership, and IT departments

THE HEALTHCARE STAFFING CRISIS

Despite having access to EHR data (like Epic) and workforce data (like UKG or Workday), most healthcare systems still rely on **manual scheduling spreadsheets, flat file imports,** or **limited rule-based software** that fails to scale with complexity.

This disconnect between **predicted demand and real-world staff allocation** leads to coverage mismatches, missed care targets, and reduced job satisfaction. Solving this requires a shift from traditional scheduling tools to **intelligent systems** that can factor in rules, constraints, and personal preferences—automatically and at scale.

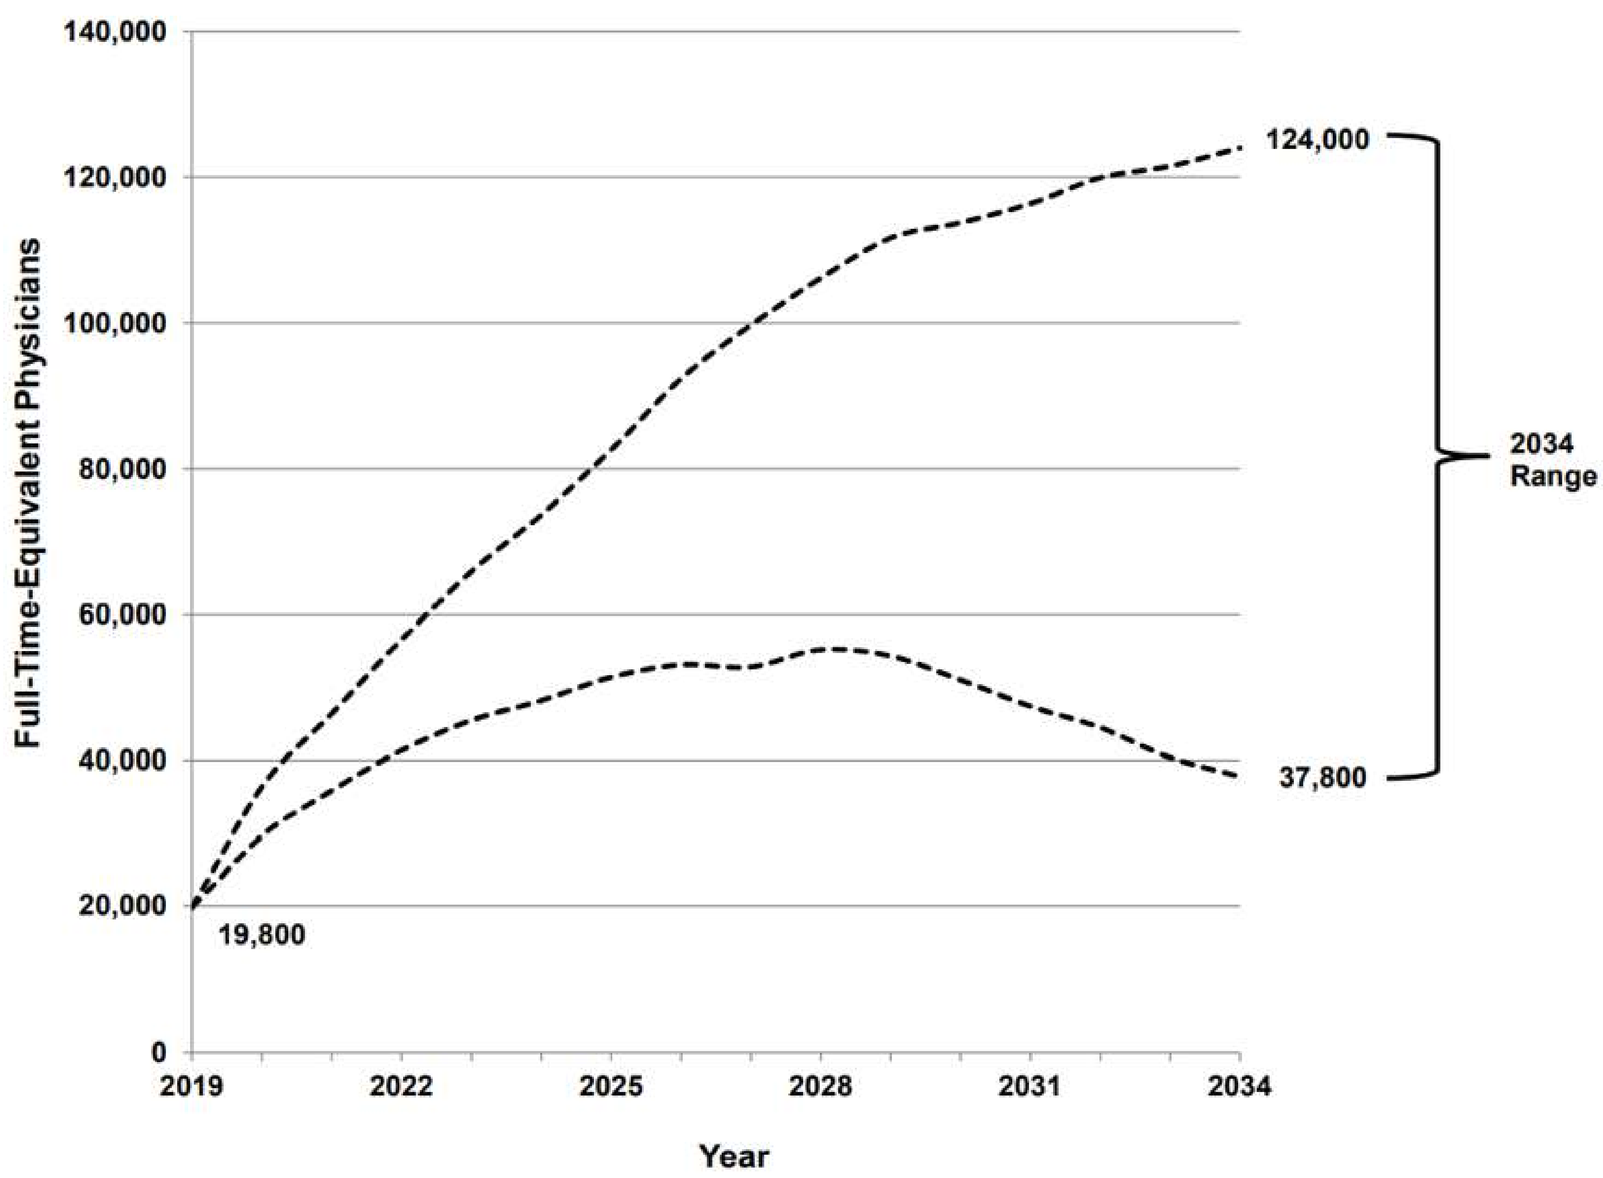

Source: Association of American Medical Colleges (AAMC)

THE CHALLENGE: COMPLEX STAFFING OPTIMIZATION

Staff scheduling in healthcare isn't just about filling shifts—it's about balancing patient demand, clinician availability, legal rules, and individual preferences, all at once. Most hospitals rely on spreadsheets or rigid software that struggle to account for:

- Varying coverage needs across units and facilities
- Staff availability and preferences
- Compliance with labor laws and union rules
- Minimizing shift swaps and manual overrides

Even with predictive models for demand and HR systems tracking availability, the gap remains: **how to intelligently match supply to demand** while respecting constraints. At scale—across hundreds of clinicians and shifts—manual scheduling breaks down.

To solve this, healthcare systems need a smarter approach that delivers **optimal assignments,** adapts in real-time, and clearly explains scheduling decisions. This is where **Linear Programming (LP)** and **AI agents** offer a powerful, scalable solution.

SOLUTION OVERVIEW: AGENTIC AI + LINEAR PROGRAMMING

To solve the complex, constraint-heavy nature of clinician scheduling, we combine Linear Programming for mathematical precision with AI agents for usability, flexibility, and explainability. This hybrid approach scales across large systems while keeping the experience intuitive for end users.

Key Components

- Linear Programming engine solves for optimal staff-to-shift allocations under real-world constraints
- Agentic AI layers provide intelligent interaction, explanations, and adaptability to shifting conditions
- Modular design integrates with Epic, Workday, and UKG—either via APIs or flat files

SOLUTION OVERVIEW: AGENTIC AI + LINEAR PROGRAMMING

Inputs & Data Flow

Input Type	Source System	Details
Demand Forecast	Epic EHR (FHIR API / Flat File)	90-day predictive staffing per unit
Clinician Supply	Workday / UKG	Availability, roles, credentials
Staff Preferences	Internal portal / Survey agent	Shift types, off-days, work-life constraints
Scheduling Rules	HR / Union policies	Fatigue limits, max shifts, mandatory rest

SOLUTION OVERVIEW:
AGENTIC AI + LINEAR PROGRAMMING

Optimization Logic (Driven by LP)

- Ensures all required coverage is met for each unit/shift
- Applies constraints like max hours/week, rest periods, skill mix
- Balances soft constraints like preferences and fairness
- Objective: Maximize shift satisfaction while minimizing conflicts and manual overrides

Hospital Staff Scheduling Optimization Architecture

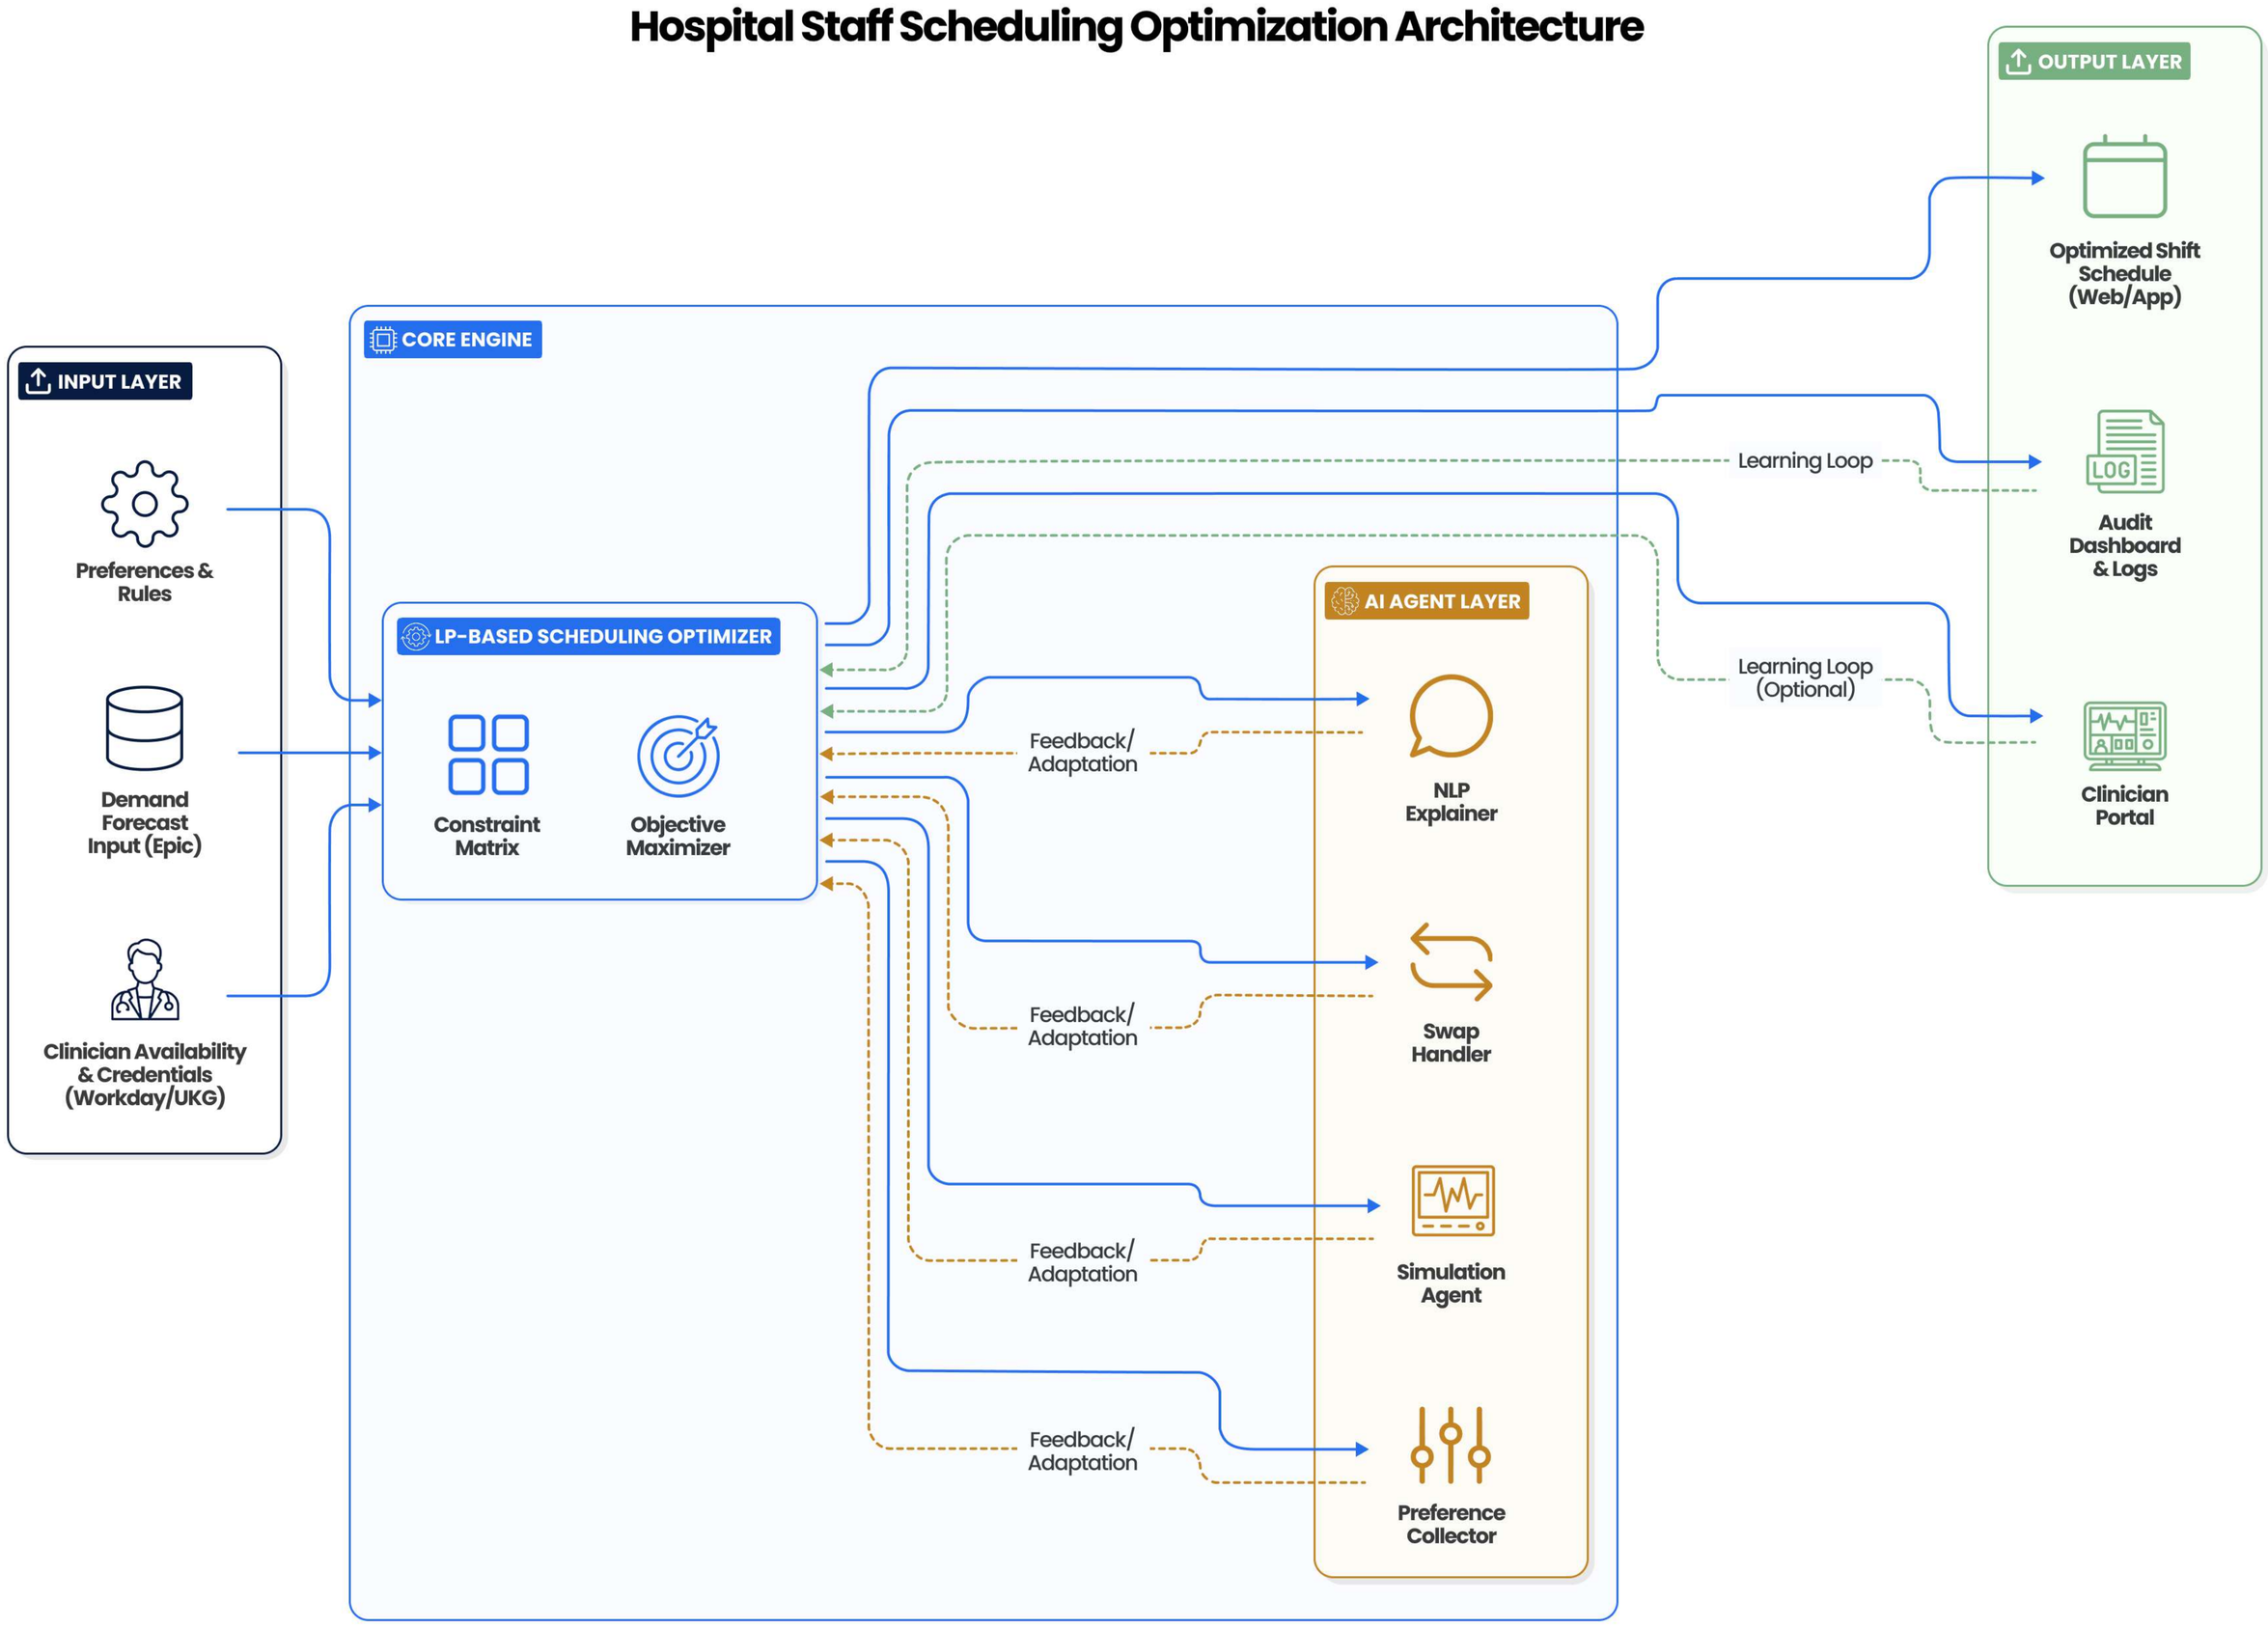

SOLUTION OVERVIEW: AGENTIC AI + LINEAR PROGRAMMING

Real-World Adaptability

- System automatically flags high-risk gaps if clinician preferences create shortfalls
- Learns from historical swaps and absences to reprioritize constraints over time
- Transparent logs help HR and clinical leads audit how and why schedule decisions were made

AI AGENT DESIGN

To make the LP-based scheduler usable, adaptive, and transparent, we embed a layer of **agentic AI components** that interact with users across roles—clinicians, schedulers, HR, and managers. These agents are purpose-built to handle domain-specific tasks, drive user engagement, and bridge the technical complexity of the optimization engine.

AI AGENT DESIGN

Key Agent Modules

Agent Module	Primary User	Function	Example Use Case
NLP Explainer	Schedulers	Interprets scheduling logic and constraint impact in natural language	“Why wasn’t Dr. Kim assigned to ICU on Tuesday?” → “Not ICU-certified; has a hard constraint on Tuesdays.”
Simulation Bot	Managers & Planners	Runs “what-if” scenarios to analyze impact of staffing changes or demand fluctuations	“What if 2 ICU nurses are unavailable next week?”
Preference Collector	Clinicians	Gathers and updates shift preferences, availability, and soft constraints via portal/ mobile interface	“I prefer morning shifts and no weekends.”
Swap Validator	Clinicians + HR	Validates shift swap requests against the LP model and institutional rules	“Can Alex swap shifts with Priya? Suggests conflict or confirms feasibility.”

AI AGENT DESIGN

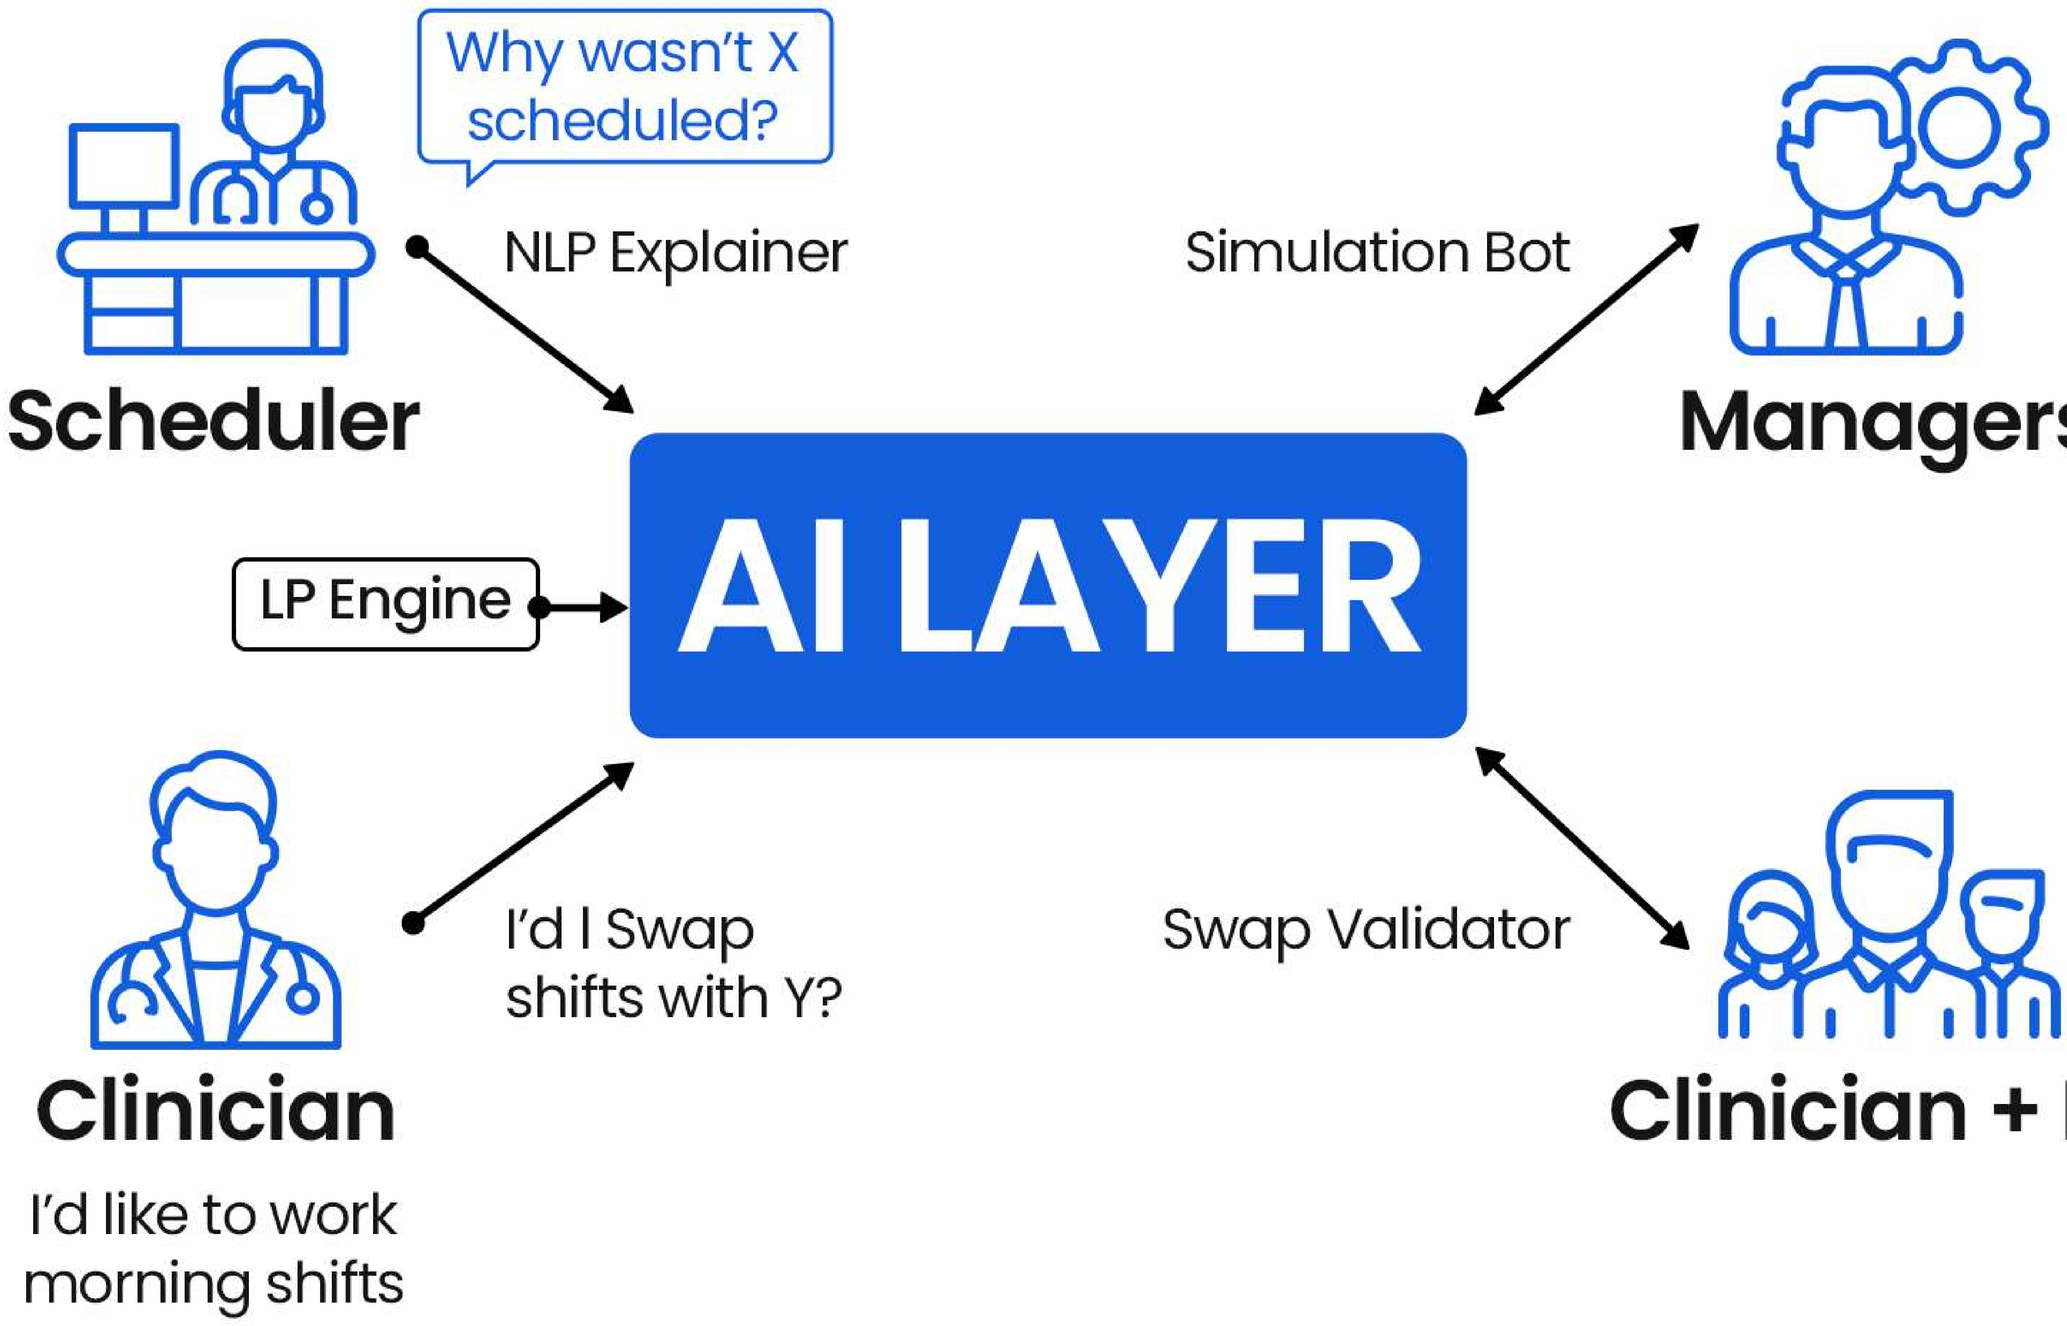

IMPLEMENTATION PHASE

Implementation Phases

We follow a five-phase rollout plan that reduces risk and delivers measurable results quickly. The goal is to validate the scheduling engine in real-world settings, starting with limited deployments and scaling gradually.

Phased Rollout Plan

Phase	Description	Deliverables	Duration
Discovery	Stakeholder interviews, constraint gathering, data source audit	Scheduling rulebook, integration map	1–2 weeks
Modeling	LP constraint model setup, input mapping, solver prototyping	Optimizer model with test data	2–3 weeks
Prototype	Integration with sample Epic/Workday/HR data, pilot agent deployment	Working demo for 1–2 units	3–4 weeks
Pilot	Real-world deployment in selected department/ unit with live feedback	Weekly schedules, preference capture, metrics	4–6 weeks
Rollout	System-wide adoption across departments, integration into clinician portals	Full coverage, ongoing constraint refinement	6–12 weeks

SECURITY, COMPLIANCE & GOVERNANCE

Our architecture is designed with security and healthcare compliance at its core—ensuring your data, workflows, and audit requirements are fully protected and transparent.

Key Safeguards

- **HIPAA Compliance**
  - All data at rest and in transit is encrypted (AES-256, TLS 1.2+).
  - Access controls are enforced with strict role-based permissions.
- **Audit Logging**
  - Every schedule change, preference update, and swap decision is logged for traceability.
  - Immutable logs available to HR and compliance leads.
- **FHIR-Native Architecture**
  - Data from Epic EHR is consumed via secure FHIR APIs.
  - Only minimal necessary data is stored for scheduling optimization.

SECURITY, COMPLIANCE & GOVERNANCE

FHIR-based Data Flow From EPIC to Scheduling Engine

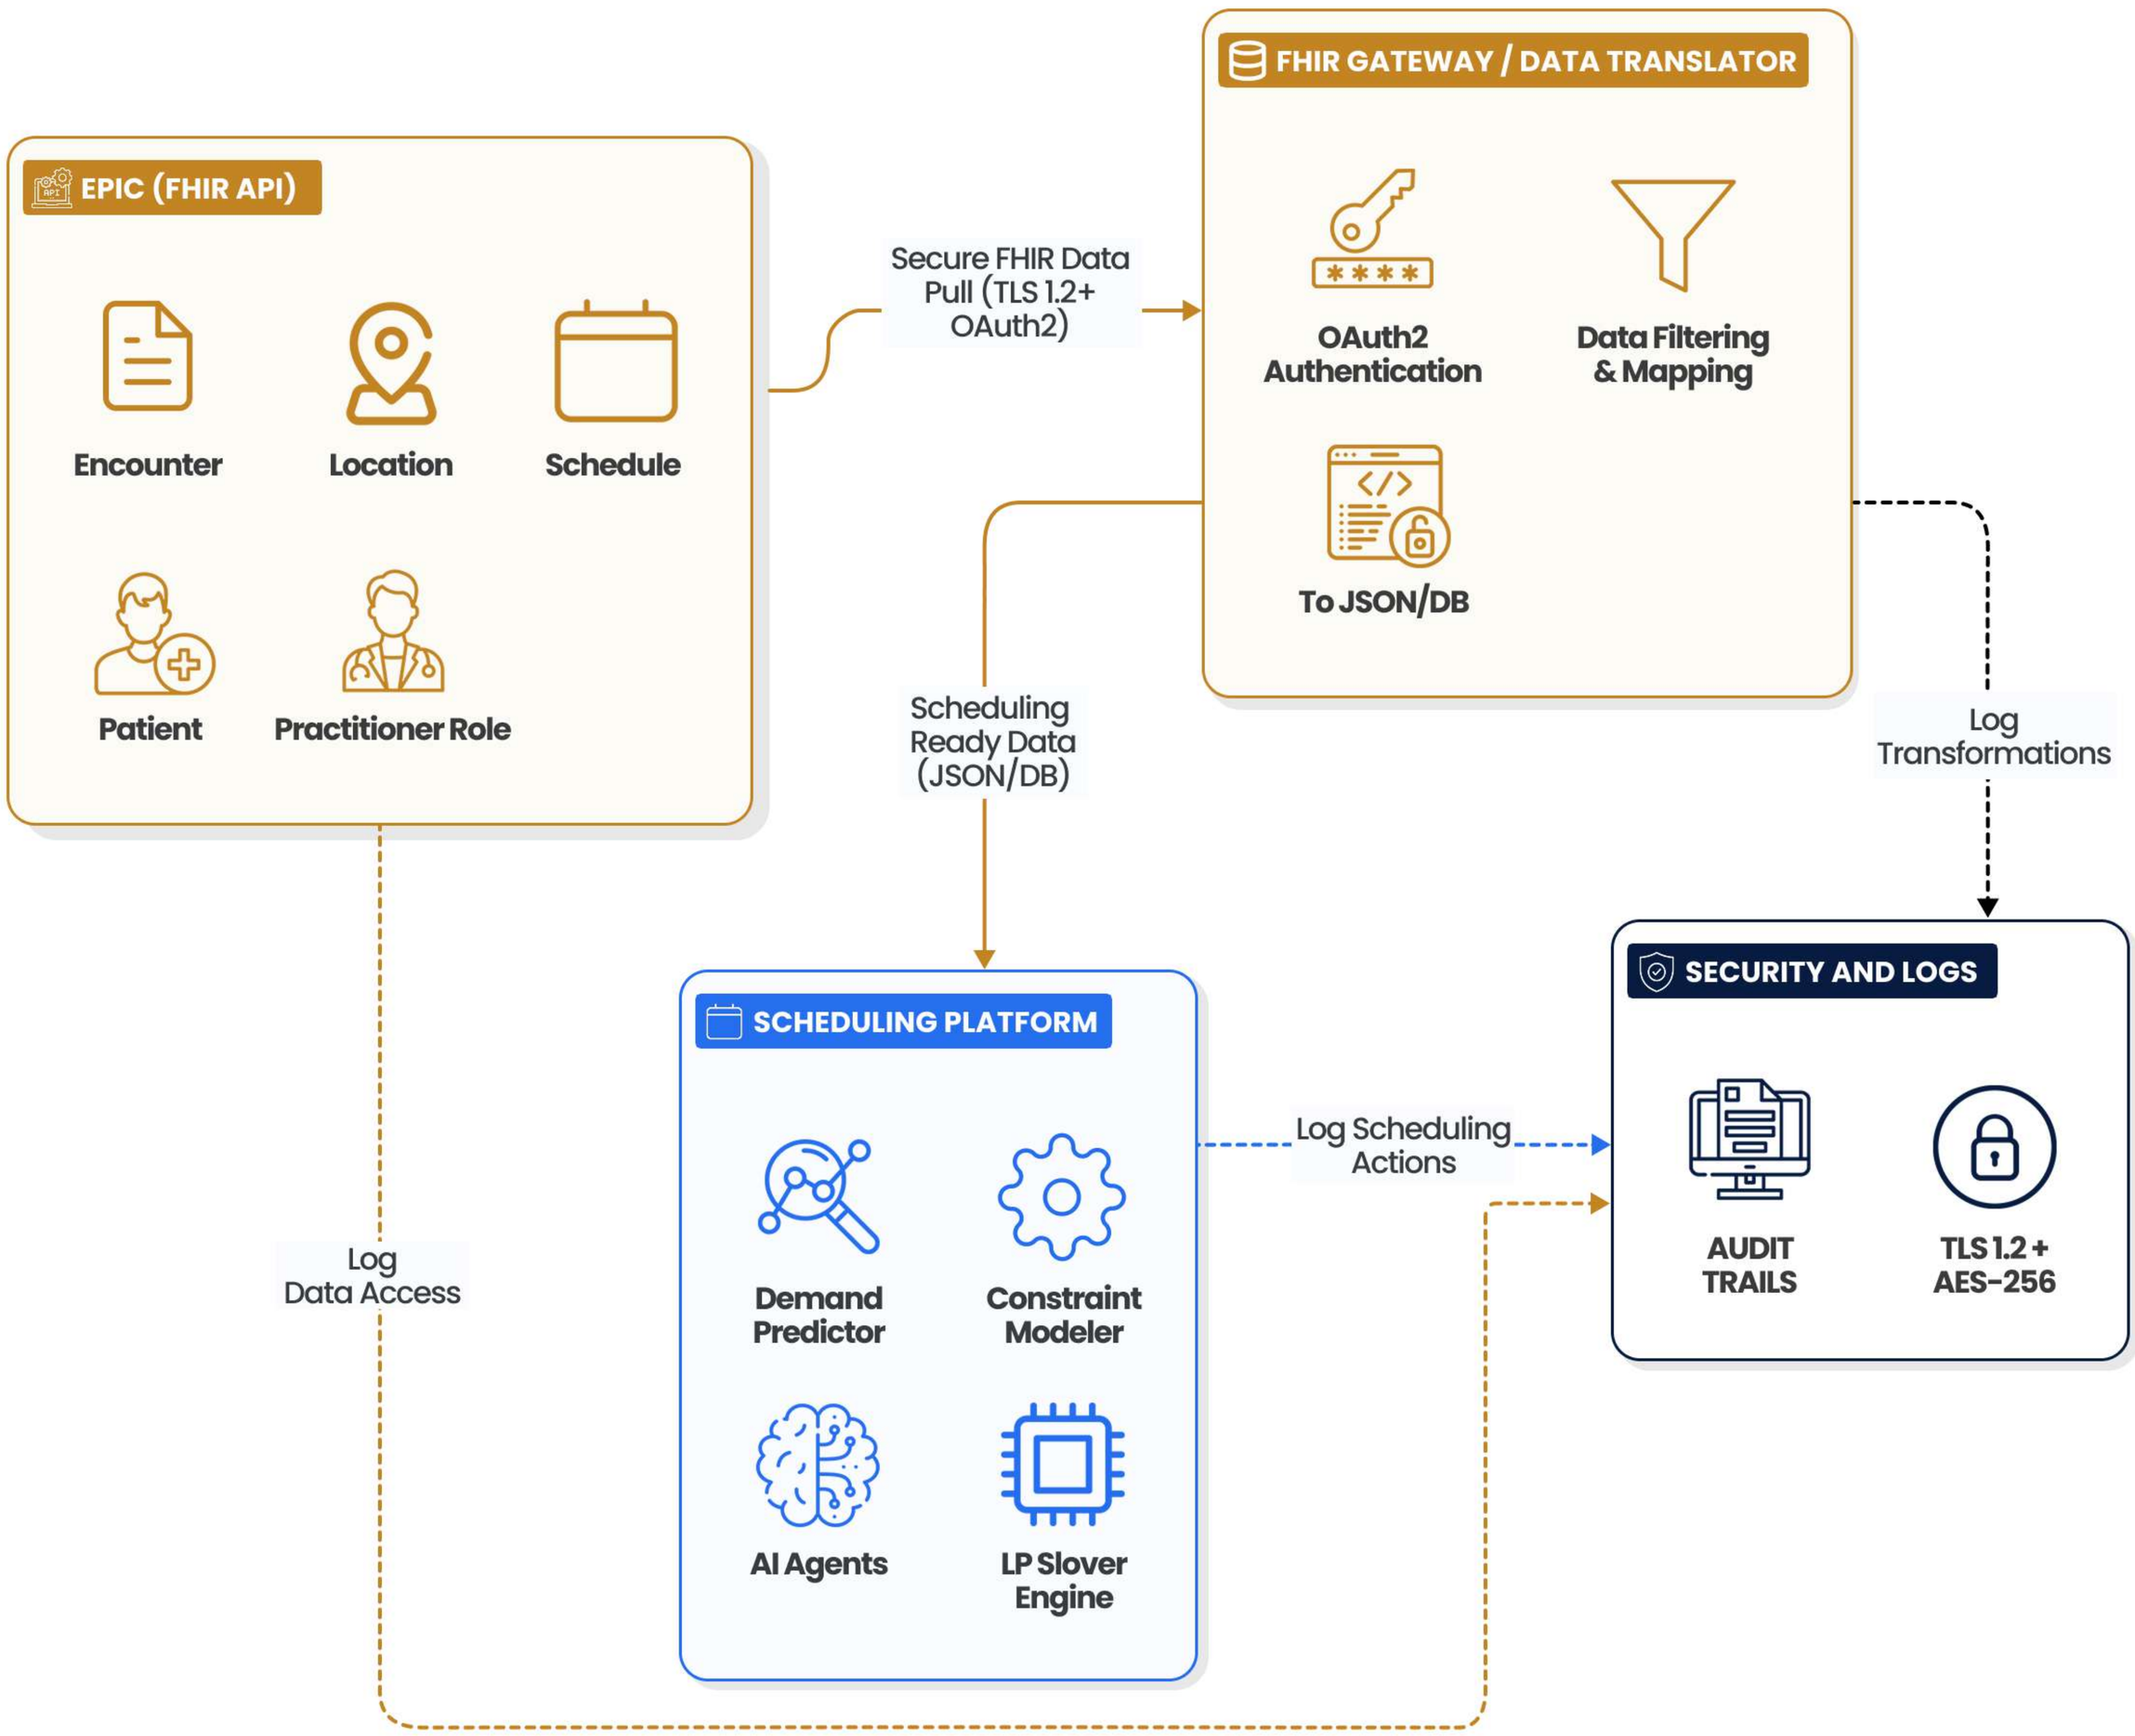

WHY WE'RE A FIT

OSP Labs is uniquely positioned to execute this solution with precision, scale, and security.

Proven Capabilities

- **FHIR, Epic, and Workday Integration Expertise**
  - Over 30 Epic integration projects and custom Workday automation builds completed.
- **Constraint-Based Optimization at Scale**
  - In-house teams with expertise in Google OR-Tools, Pyomo, and Gurobi
- **Purpose-Built AI Agents for Healthcare Use Cases**
  - NLP explainers, schedule simulators, and swap validators in production use

OSP

We are a leading software development company aiming to empower, and inspire the world with next-gen solutions. We help in simplifing every step of the development process, from system architecture design to quality delivery. Our intelligent processes enable quick deployment of enterprise-grade solutions against the toughest, and most complex challenges.

We are re-imagining how technology can empower BFSI organizations to build solutions for every day use in business applications.

With 10+ years of experience, and 200+ customers worldwide, we're leveraging technology to build the future today.

---

www.osplabs.com

Discuss Your Project

solutions@osplabs.com

Texas | California | Maryland | Mumbai