RECOVERY NETWORK
when anyone, anywhere, reaches out for help

Sovereign Intelligence Integration

CIO / CTO / VP of Engineering

How it connects, how it runs, and what your team never has to build.

Recovery Network is an API-first intelligence layer that sits above your existing EHR without replacing it. Your engineering team does not write integration code — we provide and operate the full connection and orchestration layer. Clinical data is ingested, processed, and routed through a serverless architecture using Google Cloud Functions and Pub/Sub queues. Six specialized AI systems process de-identified clinical signals in parallel, reducing signal-to-action latency to twenty-four milliseconds. Deployment is measured in hours, not months.

24ms	<24 Hours	Zero	99.9%
Signal to Council	Go-Live SLA	EHR Code Changes	Platform Uptime

HOW THE INTEGRATION WORKS

Recovery Network connects via SMART on FHIR OAuth 2.0 for Epic and Oracle Cerner. Your administrator performs a one-time OAuth handshake in the EHR app console — identical to authorizing any other health technology partner. For Kipu, Netsmart, athenahealth, and NextGen we use native REST APIs with API key authentication. For facilities without a Tier 1 EHR, our Digital Dropzone accepts HTTPS-uploaded PDFs and CSVs processed by the AI Council via OCR in under sixty seconds. No client-side code required. We own and operate the full connection layer.

SYSTEM ARCHITECTURE — EXECUTION AND DATA LAYER

Execution: Google Cloud Functions handle FHIR ingestion and rollup processing. Pub/Sub queues provide asynchronous orchestration across services. Data: each facility receives a dedicated Firestore instance — core collections are patients, encounters, signals, risk_scores, interventions, and audit_logs. PHI never leaves the tenant boundary. AI Orchestration: a multi-model pipeline processes de-identified signals via queue-based execution and consensus scoring — no raw PHI enters the AI layer. Observability: structured logging via Google Cloud Logging, Cloud Monitoring dashboards, threshold-based alerting, and immutable per-tenant audit logs. Failed events have full replay capability.

PHI BOUNDARY AND DATA FLOW

The Rollup Cloud Function is the single PHI crossing point in the architecture — the only boundary where identifiable data is transformed. No raw PHI is exposed to the AI orchestration layer. Raw EHR data flows into your dedicated GCP project. De-identified, facility-scoped tokens are passed to the AI Council. Scored recommendations are written back to your Firestore instance and surfaced through the seven clinical portals. Your security team can inspect every crossing point before go-live. Full architecture diagrams are included in the onboarding package.

RELIABILITY, DEPLOYMENT, AND OPERATIONS

The Phoenix self-healing subsystem provides automated fault detection across five infrastructure components, maintaining 99.9% uptime per calendar month. If the EHR integration goes offline, the platform continues from previously ingested data — patient risk scores do not drop to zero. Scheduled maintenance is communicated 48 hours in advance. Go-live takes under twenty-four hours from completed activation. No SDK, no developer portal, no code your team writes. Platform upgrades deploy automatically — your team is never responsible for version management.

Architecture is destiny — we have built yours for the next decade of recovery

RECOVERY NETWORK
when anyone, anywhere, reaches out for help

SYSTEM ARCHITECTURE OVERVIEW

Data flow, execution boundaries, and PHI isolation

206.353.8771 • info@recoverynetworkinc.com • recoverynetworkinc.com