TECH SOLUTIONS

AI Healthcare Governance System Architecture

PHASE - A

Approved By:

Samir Khan

Founder And Chairman of,
Classical Hunter Tech Solutions LTD
Mail: classicalhunter1@gmail.com
Phone: +880 1791-433064

Phone/WhatsApp: +880 1791-433064

Table of Contents

Executive Vision

The ***AI Healthcare Governance*** initiative is a governance-first clinical safety pilot designed to validate whether healthcare decision systems can remain explainable, auditable, and institutionally trustworthy under operational pressure.

Rather than focusing on feature-heavy hospital automation, this pilot focuses on one critical institutional question:

"Can clinical decision systems safely assist medical professionals while maintaining complete governance, traceability, and accountability?"

The Phase A pilot introduces a controlled Medication Interaction & Allergy Checker capable of:

- Detecting dangerous drug interactions
- Identifying allergy conflicts
- Classifying clinical risk
- Enforcing controlled override procedures
- Maintaining immutable audit history
- Supporting explainable AI-assisted decisions

The system architecture prioritizes:

- Human-supervised decision making
- Controlled escalation workflows
- Deterministic safety enforcement
- Regulatory-grade auditability
- AI-assisted governance support

This initiative establishes a foundational governance infrastructure upon which future healthcare intelligence systems may safely evolve.

Clinical Scenario

Example Institutional Workflow:

Patient Admission

A patient arrives at a healthcare facility with a known Penicillin allergy. The attending physician performs an evaluation and enters a prescription into the system.

Step 1 — Prescription Submission

The physician submits:

- Patient information
- Allergy history
- Medication recommendation
- Dosage instructions

The system immediately validates:

- Existing allergies
- Drug-to-drug interactions
- Risk severity
- Patient safety constraints

Step 2 — Conflict Detection

The AI-assisted governance engine identifies a conflict:

Detection Type	Result
Drug-Allergy Conflict	Penicillin Allergy Detected
Risk Level	HIGH RISK
Clinical Severity	Critical
Action Required	Immediate Restriction

The system prevents unsafe continuation automatically.

Step 3 — Risk Classification

The system categorizes the event using governance rules and historical validation data.

Risk Evaluation Factors

- Known allergy severity
- Medication compatibility
- Previous patient reactions
- Institutional risk policy
- Clinical confidence score

Step 4 — Constraint Enforcement

Because the prescription exceeds permitted risk thresholds:

- Prescription execution is blocked
- Override authority becomes restricted
- Supervisor review is required
- Clinical justification becomes mandatory

The physician cannot bypass the restriction independently.

Step 5 — Explainability Layer

The system generates an explainable decision summary:

Governance Field	Example
Triggered Rule	Penicillin-Allergy-Conflict
Severity Level	High
Decision Source	AI + Rule Validation
Timestamp	Logged
Reviewing Authority	Required

Every decision remains transparent and reviewable.

Step 6 — Controlled Override Workflow

If exceptional clinical circumstances exist:

1. Supervisor receives escalation request
2. Justification is reviewed
3. Override approval or rejection is recorded
4. All participants are logged
5. Final outcome becomes immutable

This ensures governance remains human-controlled.

Step 7 — Immutable Audit Trail

Every event becomes permanently traceable:

- Doctor identity
- Supervisor approval
- Timestamp
- Risk evaluation
- Rule version

- Final outcome
- System-generated reasoning

No audit record can be modified retroactively.

Step 8 — Regulator Scenario

Years later, a regulator requests the full clinical decision history for the patient encounter.

The system can provide:

- Complete timeline reconstruction
- Cryptographically verifiable audit integrity
- Decision explainability
- Override accountability
- Governance compliance evidence

This validates institutional trustworthiness.

Clinical Decision Flow Chart

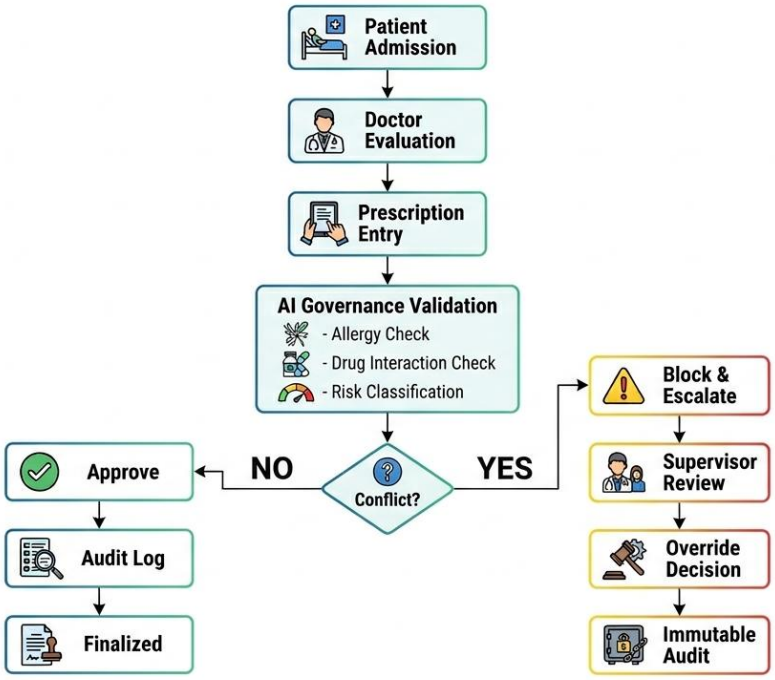

fig: Visual representation of system workflow

Governance Principles

Core Institutional Principles:

1. **Immutable Audit Architecture:** Every critical action is permanently logged and cryptographically verifiable.
2. **Explainable Decisions:** No black-box decision execution is permitted. All outcomes remain reviewable and understandable.
3. **Human Override Authority:** AI may assist decision-making, but final authority remains under controlled human governance.
4. **Constraint-First Safety:** Unsafe actions are prevented before execution rather than reviewed afterward.

fig: Visual representation of governance

5. **Traceable Escalation:** All escalation paths, approvals, and interventions remain fully attributable.
6. **Governance Before AI:** AI capability expansion is permitted only after governance reliability is validated.

System Architecture

Core Architectural Layers:

Layer	Responsibility
Client Layer	Doctor, Supervisor, Auditor Interfaces
Access Layer	Authentication & Role Governance
Application Layer	Prescription & Override Services
Governance Layer	AI + Rule Validation Engine
Audit Layer	Immutable Event Logging
Storage Layer	PostgreSQL 16 Data Infrastructure

High-Level Architecture:

The platform follows a layered governance-oriented architecture designed for controlled healthcare decision support.

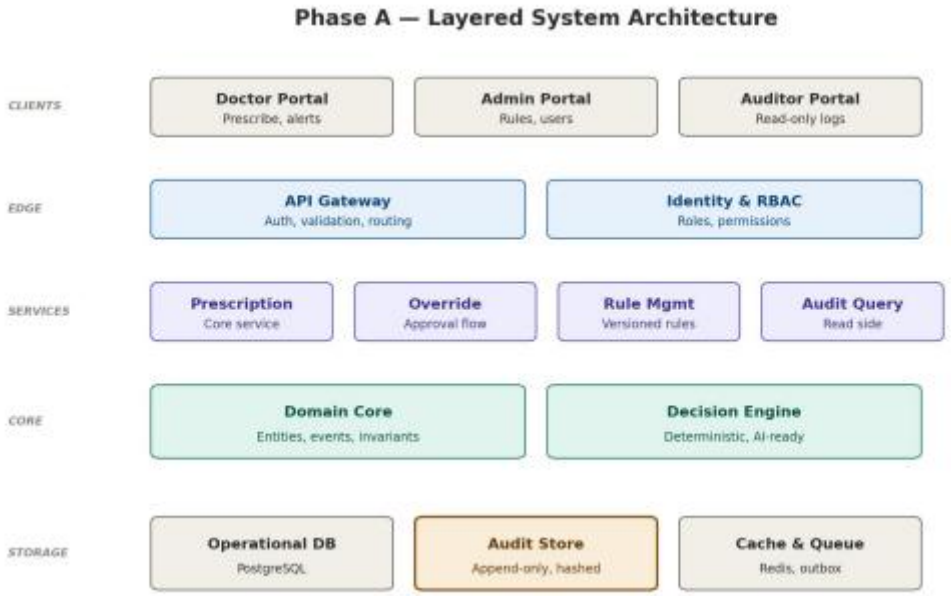

fig: Diagram of system architecture.

Technology Stack:

The technology foundation for Phase A has been selected based on institutional reliability, governance suitability, ecosystem maturity, and long-term scalability.

The objective was not to maximize technological complexity, but to establish a stable, auditable, and operationally trustworthy healthcare governance environment.

Layer	Technology	Strategic Purpose
Backend Platform	Laravel 13 (PHP 8.3+)	Mature development ecosystem with strong security, governance flexibility, rapid implementation capability, and structured RBAC integration
Database Infrastructure	PostgreSQL 16	Enterprise-grade relational database designed for transactional consistency, audit reliability, and healthcare-compatible data integrity
Cache & Queue Services	Redis 7	High-performance caching, session management, asynchronous processing, and transactional event coordination

Frontend Framework	Tailwind CSS + Flowbite + Alpine.js	Lightweight, governance-oriented interface delivery with fast rendering, minimal operational complexity, and efficient Laravel integration
Service Architecture	Containerized Microservices	Clear governance boundaries, independent deployment capability, operational isolation, and scalable service ownership
Authentication & Access Control	Sanctum (Phase A) with future Keycloak integration pathway	Secure role-based authentication with future enterprise identity federation readiness
Observability & Monitoring	OpenTelemetry + ELK Stack	Centralized operational monitoring, system visibility, audit-supportive log aggregation, and traceability management
Deployment Environment	Docker-Based Container Infrastructure	Consistent deployment lifecycle, controlled scalability, operational portability, and infrastructure standardization

Audit & Security

Regulatory-Grade Auditability

The audit subsystem serves as the institutional trust foundation of the platform.

Audit Philosophy

The system assumes that every critical clinical action may eventually require external institutional review.

Therefore:

- Every action is attributable
- Every decision is explainable
- Every override is governed
- Every event is traceable

Security & Governance Features

Feature	Purpose
Immutable Audit Logging	Prevents retroactive modification
Cryptographic Verification	Detects tampering attempts
Role-Based Access Control	Restricts unauthorized actions

Explainable Decision Tracking	Maintains transparency
Controlled Override Workflow	Enforces governance escalation
Encrypted Patient Data	Protects sensitive information

Validation Strategy

Governance Validation Approach

The pilot focuses on validating institutional reliability before operational scaling.

Validation Areas

- **Shadow-Mode Clinical Testing:** The system operates alongside manual review to validate safety consistency.
- **Adversarial Workflow Testing:** High-risk and edge-case scenarios are intentionally simulated.
- **Audit Verification:** Audit chain integrity and traceability are continuously verified.
- **Override Governance Testing:** Supervisor escalation workflows are stress-tested under operational pressure.
- **Controlled Rollout:** Deployment occurs gradually to ensure governance stability before expansion.

Delivery Roadmap

Phase A Delivery Structure:

Phase	Focus Area
Foundation	Authentication, RBAC, Audit Infrastructure, Core Validation Engine
Governance	Override Workflow, Explainability, Risk Escalation
Validation	Security Testing, Hardening, Audit Verification, Pilot Deployment

Estimated Delivery Duration:

Approximately ***16 weeks*** after project start.

Delivery Team Structure:

Role	Allocation
Project Manager	1
UI/UX Designer	1
Frontend Developer	1
Senior Software Developer	1
Software Developers	2
Clinical Advisor	1
QA Tester	1

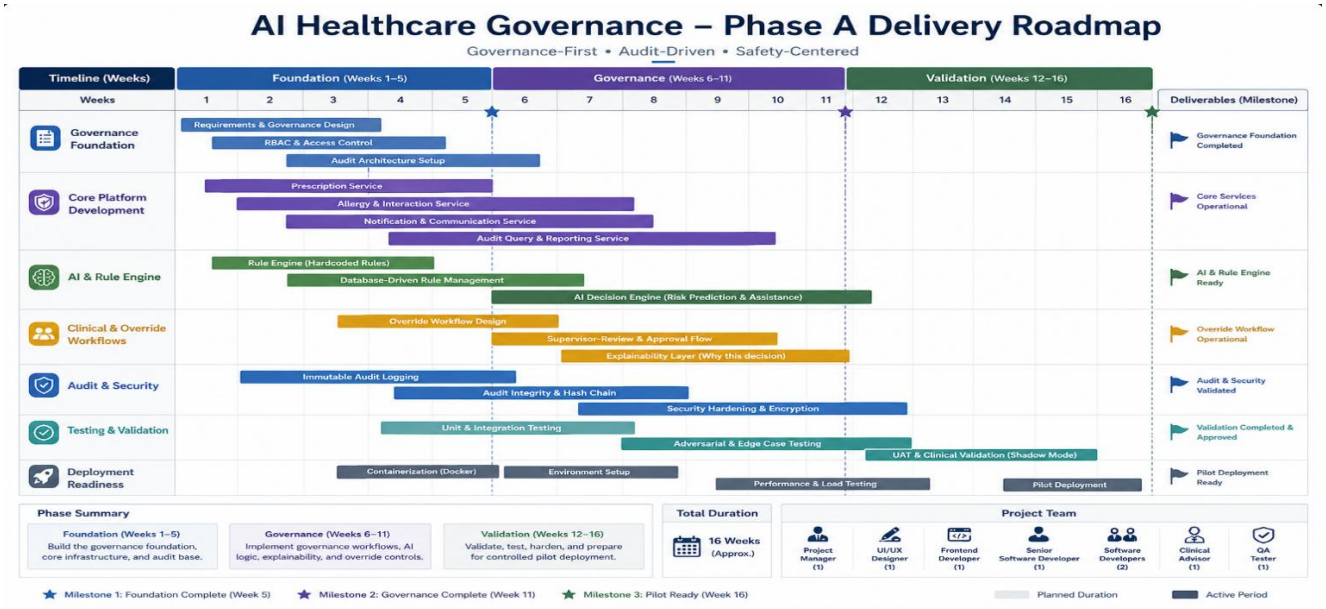

fig: <u>Gantt chart of delivery roadmap.</u>

Strategic Expansion Path

The current pilot intentionally limits operational scope in order to validate governance reliability first.

Future expansion possibilities may include:

- AI-assisted triage systems
- Clinical recommendation intelligence
- Hospital resource coordination
- Adaptive governance models
- Predictive patient safety systems

However:

"Expansion is permitted only after governance validation demonstrates institutional reliability and operational trustworthiness."

Interactive Visual Representation

Visual Workflow Demonstration

For a complete visual presentation of the governance workflow, clinical interaction flow, and operational interface structure, please refer to the secured presentation environment below.

This environment is intended exclusively for workflow visualization and institutional process demonstration purposes.

URL : https://samir-trial.hostingprovider.site/

Email : admin@samir-trial.local

Password: SamirTrial-FQ7lg7Bcl0ndIYg7ou

Note: This environment is provided solely as a workflow and governance presentation platform. It is not intended for production clinical use.

About CHTS

Classical Hunter Tech Solutions LTD (CHTS) is a full-service IT company specializing in:

- *AI (Artificial Intelligent) Automation System Development*
- *Mobile Application Development*
- *Custom Software Development*
- *School and Hospital Management Solutions*
- *Web Design and Development*
- *UI/UX Design*
- *Intercom Network Solutions*
- *Network Architecture and Maintenance*
- *E-commerce and ERP Solutions*
- *Portfolio and Blog Site Development*
- *All kinds of tailored IT solutions for businesses and institutions*

With years of experience and a skilled team, we are committed to delivering innovative and user-friendly digital

systems to transform the way our clients operate.

Contact information:

Email : classicalhunter1@gmail.com

Phone : +880 1791-433064

Linked List : linkedin.com/company/classical-hunter-tech-solutions

Conclusion

Phase A is not designed to maximize features.

It is designed to validate whether healthcare decision systems can remain:

- Governable

- Explainable
- Auditable
- Safe
- Institutionally trustworthy

under real operational conditions.

The long-term value of this initiative is not merely technological capability, but the establishment of a healthcare governance foundation capable of supporting future intelligent clinical systems responsibly.

Success will ultimately be measured by one institutional question:

"Can every critical healthcare decision remain explainable, reviewable, and cryptographically verifiable years after execution?"

If the answer remains yes, the governance architecture succeeds.