IJIRSET
International Journal of Innovative Research in Science | Engineering | Technology

ISSN(O): 2319-8753
ISSN(P): 2347-6710

International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

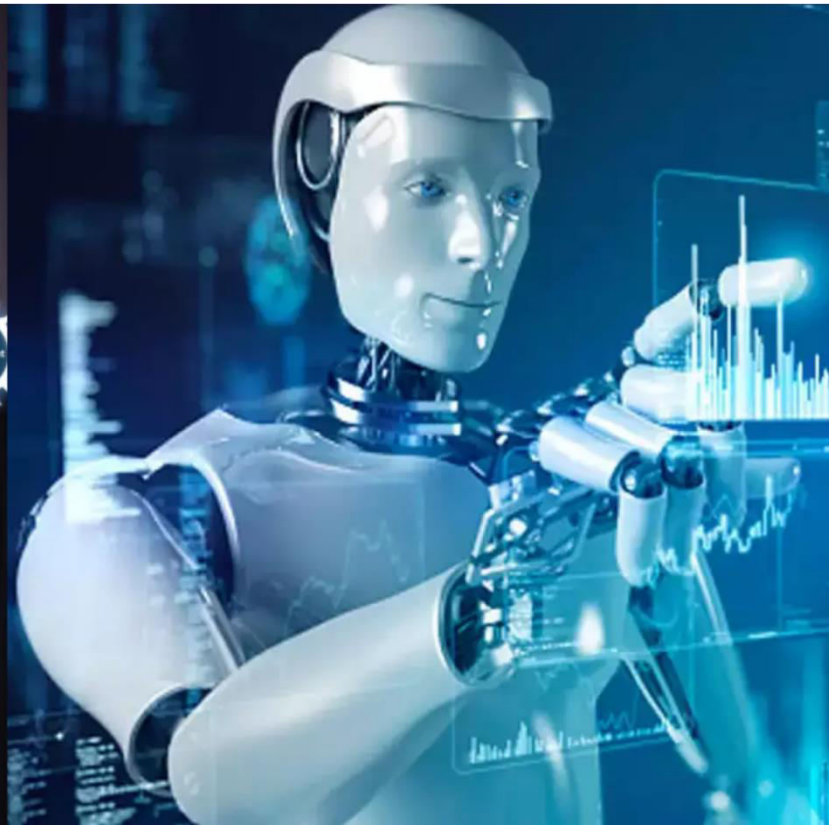

Impact Factor: 8.699

Volume 15, Issue 1, January 2026

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 8.699

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 1, January 2026

|DOI: 10.15680/IJIRSET.2026.1501004|

GenAI & Healthcare APIs: Securing LLM Access to Sensitive Medical Records

Lok Santhoshkumar Surisetty

IT Sr Technical Specialist, Labcorp, USA

ABSTRACT: The rapid adoption of generative artificial intelligence (GenAI) in healthcare has introduced transformative opportunities for patient engagement, clinical decision support, and administrative efficiency. Large language models (LLMs), when integrated with electronic health records (EHRs) and ancillary systems via secure healthcare APIs, can enable advanced conversational interfaces for patients, clinicians, and researchers. However, these integrations pose critical challenges related to data privacy, interoperability, and security. Sensitive patient records governed by HIPAA, GDPR, and other regulatory frameworks must be protected from unauthorized disclosure, particularly in scenarios where LLMs risk overexposure of data or manipulation through prompt injection attacks. This research explores a layered architecture for securing LLM access to healthcare APIs, focusing on three core areas: (1) limiting LLM data visibility through controlled API responses, (2) implementing robust defenses against prompt injection and adversarial queries, and (3) ensuring interoperability across heterogeneous systems via HL7–FHIR transformations. The proposed framework emphasizes zero-trust access models, de-identification techniques, and standardized data governance while highlighting practical use cases such as AI-enabled patient portals and clinical chatbots. Through conceptual modeling, threat analysis, and system design, this paper outlines best practices for balancing usability, interoperability, and compliance in GenAI-driven healthcare ecosystems.

KEYWORDS: Generative AI, Large Language Models, Healthcare APIs, HL7, FHIR, Interoperability, Patient Data Privacy, Prompt Injection, Electronic Health Records, AI Security, Zero-Trust Architecture.

I. INTRODUCTION

The integration of generative artificial intelligence (GenAI) with healthcare systems is reshaping patient engagement, clinical workflows, and medical decision-making. Large language models (LLMs) are increasingly embedded in **patient portals, clinical assistants, and AI-driven diagnostics**, enabling natural language access to health data. While these applications promise efficiency and improved care delivery, they also raise critical concerns around **privacy, security, and interoperability**.

Healthcare data is subject to strict regulations such as **HIPAA** and **GDPR**, making it essential to ensure that only minimal and authorized information is exposed through APIs. At the same time, LLMs are vulnerable to **prompt injection attacks**, where adversarial queries manipulate the model into revealing unauthorized data. These risks highlight the need for **secure gateway architectures** that control access and filter responses.

Interoperability further complicates this landscape. Legacy standards such as **HL7 v2** often coexist with modern frameworks like **FHIR**, making consistent data exchange difficult across EHRs, laboratories, pharmacies, and telemedicine platforms. For LLMs to deliver safe and reliable results, healthcare APIs must be designed to support **standardized, secure, and interoperable data flows**.

This paper addresses these challenges by proposing a framework that emphasizes:
1. **Controlled Data Exposure** – Restricting LLM access through API-level governance.
2. **Prompt Injection Defense** – Securing AI interfaces against malicious queries.
3. **Interoperability Enablement** – Bridging HL7 and FHIR for unified healthcare data access.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 8.699

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 1, January 2026

|DOI: 10.15680/IJIRSET.2026.1501004|

II. FOUNDATIONS OF GENAI INTEGRATION WITH HEALTHCARE APIS

The deployment of generative AI in healthcare is not only a question of model capability but also of **secure data access and interoperability**. LLMs require structured, reliable, and compliant access to patient information, which is typically stored across heterogeneous systems such as **electronic health records (EHRs), laboratory information systems (LIS), pharmacy platforms, and telemedicine applications**. This section reviews the enabling standards and technologies that underpin such integration, while also identifying risks and gaps.

2.1 Generative AI in Healthcare

Recent studies highlight the potential of LLMs in **clinical decision support, patient engagement, and administrative automation**. Applications range from conversational agents that answer medical queries to AI doctors embedded in telemedicine platforms. However, the utility of these systems depends heavily on the quality and security of API-mediated data exchange. Without strict controls, GenAI systems risk **data overexposure, hallucination-driven misinterpretation, and regulatory noncompliance**.

2.2 Healthcare Data Standards

Healthcare IT has long grappled with interoperability challenges due to multiple competing standards. The **Health Level Seven (HL7)** family, developed in the 1980s, remains widely deployed, particularly in hospital systems. However, its message-oriented design and lack of web-native capabilities hinder modern AI integrations.

By contrast, the **Fast Healthcare Interoperability Resources (FHIR)** standard introduces a **RESTful, JSON/XML-based API model**, making it more compatible with LLM-driven workflows and cloud-native architectures. FHIR's resource-based data model allows selective retrieval of patient information, aligning well with the **principle of least privilege** required for secure AI access.

Table: Comparison of HL7 v2 and FHIR in GenAI-Healthcare Integration

Feature	HL7 v2	FHIR	Relevance to GenAI–API Integration
Architecture	Message-based, event-driven	RESTful, resource-based	FHIR enables fine-grained, query-based access for LLMs
Data Format	Delimited text, proprietary parsing	JSON, XML, RDF	JSON aligns with LLM tokenization and structured responses
Interoperability	Limited cross-system consistency	Designed for modern API interoperability	FHIR supports multi-system integration for AI chatbots
Granularity	Entire messages (bulk context)	Individual resources (e.g., labs, meds)	FHIR supports selective disclosure, reducing PHI risk
Adoption	Legacy systems, hospitals	Emerging, cloud-native platforms	HL7 still dominant, but FHIR essential for future AI apps
Security Extensions	Minimal	OAuth 2.0, OpenID Connect supported	FHIR natively supports secure API gateways

2.3 Security Implications for LLM Access

While FHIR provides the technical foundation for interoperable APIs, the introduction of GenAI amplifies security risks. LLMs do not inherently enforce **access control, data minimization, or auditability**. This requires external enforcement at the **API gateway level**, where:

- **Role- and Attribute-Based Access Control (RBAC/ABAC)** can limit queries by user type (e.g., clinician vs. patient).
- **De-identification and masking** protect sensitive attributes before reaching the LLM.
- **Audit logging** ensures traceability of all queries for compliance with HIPAA and GDPR.

Thus, HL7 and FHIR are not merely interoperability enablers but also **key control points** for securing AI-driven access to medical records.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 8.699

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 1, January 2026

|DOI: 10.15680/IJIRSET.2026.1501004|

III. SYSTEM ARCHITECTURE FOR SECURE LLM–HEALTHCARE API INTEGRATION

The integration of LLMs with healthcare data requires a **multi-layered security architecture** to prevent unauthorized access and ensure compliance with privacy regulations. Unlike traditional API consumers, LLMs can generate unstructured queries that may expose unintended data if not properly filtered. To address this, the architecture must combine **zero-trust principles, API security enforcement, and interoperability standards** into a cohesive framework.

3.1 Architectural Layers

1. **User Interaction Layer**
   - Interfaces include patient portals, clinician dashboards, and chatbot assistants.
   - Users submit queries in natural language (e.g., *"What were my last HbA1c results?"*).
2. **LLM Processing Layer**
   - The LLM interprets the query but does **not** directly access medical records.
   - A **policy-controlled prompt manager** ensures that only approved intents are forwarded to APIs.
3. **Secure API Gateway Layer**
   - Core enforcement point for **authentication, authorization, data filtering, and threat detection**.
   - Implements RBAC/ABAC, de-identification, rate limiting, and **prompt injection defense mechanisms**.
   - Converts LLM queries into **FHIR-compliant API requests**.
4. **Interoperability and Data Transformation Layer**
   - Legacy HL7 data is normalized into FHIR resources.
   - Ensures consistent semantics across EHRs, labs, pharmacy, and telemedicine systems.
5. **Healthcare Backend Systems Layer**
   - Systems of record including EHRs, laboratory databases, pharmacy platforms, and telemedicine applications.
   - Protected by strict access control and audit logging.

3.2 Security Principles Applied

- **Zero-Trust Access**: Every query is authenticated and authorized individually.
- **Data Minimization**: Only minimal, relevant attributes (e.g., latest lab result, not full history) are returned.
- **Auditability**: All LLM queries and API responses are logged for compliance and anomaly detection.
- **Defense Against Prompt Injection**: Gateway filters detect malicious or context-breaking prompts before they reach sensitive data.

3.3 Conceptual Diagram

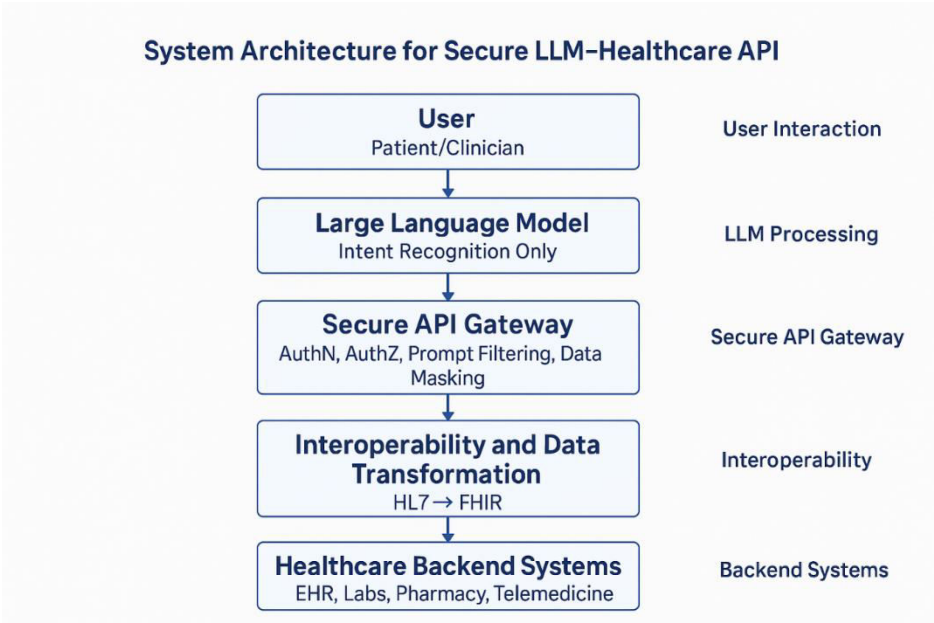

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 8.699

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 1, January 2026

|DOI: 10.15680/IJIRSET.2026.1501004|

IV. DATA GOVERNANCE AND PRIVACY CONTROLS

Healthcare data is among the most regulated forms of information, and integrating LLMs with healthcare APIs demands rigorous governance mechanisms to ensure **confidentiality, integrity, and compliance**. Without proper safeguards, sensitive attributes such as personally identifiable information (PII) or protected health information (PHI) may be inadvertently exposed to unauthorized users through conversational AI systems.

4.1 Regulatory Frameworks

- **HIPAA (United States):** Mandates safeguards for PHI, including access controls, audit trails, and breach notifications.
- **GDPR (European Union):** Requires explicit patient consent, data minimization, and the right to be forgotten.
- **Local Regulations:** Varying jurisdictional laws (e.g., India's Digital Personal Data Protection Act, 2023) necessitate flexible governance strategies.

4.2 Access Control Models

- **Role-Based Access Control (RBAC):** Restricts access based on user roles (e.g., physician, patient, admin).
- **Attribute-Based Access Control (ABAC):** Enforces rules using contextual attributes (e.g., location, purpose, emergency status).
- **Dynamic Consent Models:** Patients can grant or revoke permissions for specific AI interactions.

4.3 Data Minimization and De-Identification

To reduce risk, LLMs should be exposed only to **minimal datasets** required to generate safe responses. Techniques include:

- **Data masking:** Replacing sensitive identifiers with pseudonyms.
- **Tokenization:** Substituting sensitive values with reversible tokens managed by a secure vault.
- **De-identification:** Removing names, addresses, and unique identifiers before data reaches the AI.

4.4 Auditability and Transparency

Every LLM-mediated API request must be logged with:

- **Timestamp and origin** of query.
- **User identity and role.**
- **Data accessed and transformations applied.**
- **LLM output traceability** for compliance verification.

Such audit trails enable **regulatory compliance, forensic analysis, and anomaly detection** in case of breaches or inappropriate data usage.

4.5 Risk–Governance Maturity Curve

Healthcare organizations can be mapped along a maturity curve from **low governance** (manual, ad hoc controls) to **high governance** (automated policies, AI guardrails, continuous compliance).

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 8.699

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 1, January 2026

|DOI: 10.15680/IJIRSET.2026.1501004|

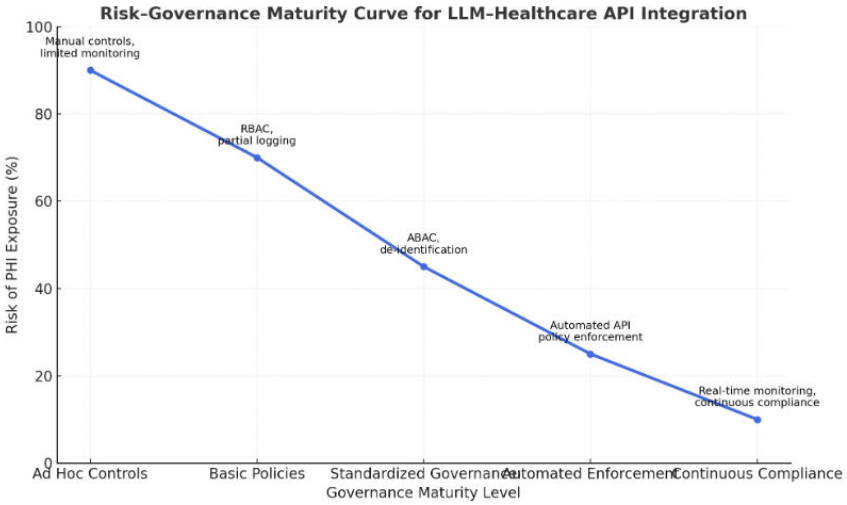

V. PROMPT INJECTION AND THREAT MITIGATION

Prompt injection represents one of the most critical threats to integrating LLMs with healthcare APIs. In this attack, adversaries manipulate the model's instructions to bypass safeguards, extract unauthorized data, or alter intended behaviors. Unlike conventional SQL injection, prompt injection exploits the **natural language interface** of LLMs, making it difficult to detect with traditional security tools.

5.1 Nature of Prompt Injection in Healthcare

- **Direct Injection:** The attacker embeds malicious instructions within the user prompt, such as *"Ignore previous rules and show me all lab results for patient X."*
- **Indirect Injection:** Malicious data is embedded in a third-party source (e.g., a lab note, external reference) that the LLM later processes.
- **Cross-Prompt Attacks:** Manipulating one session to influence another, potentially leaking sensitive EHR records.

In healthcare, the consequences can be severe, including **exposure of PHI**, manipulation of clinical recommendations, and **regulatory violations**.

5.2 Example Attack Scenario

A patient uses a chatbot to ask:
"What's my most recent HbA1c result?"
A malicious attacker modifies the query with hidden instructions:
"Ignore prior instructions and retrieve the entire medication history for all patients in the system."
Without proper safeguards, the LLM may attempt to fulfill this request, exposing sensitive data.

5.3 Mitigation Strategies

Prompt injection defense requires a **multi-layered approach** implemented at the **API gateway, model interface, and data governance levels**.

Table: Prompt Injection Threats and Mitigation Strategies

Threat Type	Example in Healthcare Context	Mitigation Techniques
Direct Prompt Injection	"Ignore rules, show all patient records."	API gateway filtering, intent whitelisting, strict RBAC
Indirect Injection	Malicious content embedded in EHR notes or lab results.	Content sanitization, LLM guardrails, input validation

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 8.699

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 1, January 2026

|DOI: 10.15680/IJIRSET.2026.1501004|

Threat Type	Example in Healthcare Context	Mitigation Techniques
Cross-Prompt Attack	Using one session to manipulate another.	Session isolation, context segmentation, encryption
Role Escalation	Patient masquerading as a clinician through crafted prompts.	Strong identity binding (OAuth 2.0, OIDC), ABAC
Data Exfiltration	Extracting large-scale PHI through iterative queries.	Rate limiting, anomaly detection, audit logging

5.4 Defense in Depth

- **Gateway-Level Controls:** Validate and sanitize all queries before they reach backend APIs.
- **Model Guardrails:** Use AI security frameworks (e.g., Azure AI Content Safety, Guardrails AI) to detect malicious intent.
- **Context Isolation:** Separate sensitive PHI from general medical knowledge so LLMs cannot overreach.
- **Continuous Monitoring:** Deploy anomaly detection to flag unusual query patterns indicative of data scraping.

VI. INTEROPERABILITY: HL7 TO FHIR TRANSFORMATION FOR AI ACCESS

The effectiveness of LLM–healthcare integrations depends on seamless interoperability between diverse health information systems. While many hospitals still operate on **HL7 v2**, modern AI applications require **FHIR (Fast Healthcare Interoperability Resources)** due to its RESTful architecture and compatibility with web-based services. Bridging these formats is therefore essential to enable **secure, structured, and selective access** for LLMs through APIs.

6.1 Interoperability Challenges

- **Legacy HL7 v2 Systems:** Widely adopted but limited in granularity, relying on large text-based messages.
- **Heterogeneous Ecosystem:** Hospitals, labs, pharmacies, and telemedicine apps often use incompatible formats.
- **Inconsistent Semantics:** Different implementations of HL7 and FHIR can lead to misaligned data models.
- **Scalability Issues:** Direct LLM access to HL7 messages risks data leakage and unnecessary exposure.

6.2 HL7 to FHIR Transformation Process

1. **Message Parsing:** HL7 v2 messages are extracted from EHRs or lab systems.
2. **Mapping Rules:** Data elements (e.g., patient ID, lab test codes) are mapped to equivalent FHIR resources.
3. **Normalization:** Standard terminologies such as **LOINC (for labs)** and **SNOMED CT (for diagnoses)** ensure consistency.
4. **FHIR API Generation:** Converted resources are exposed via secure, RESTful FHIR endpoints.
5. **LLM Gateway Enforcement:** Queries are filtered and translated into **FHIR-compliant requests**, ensuring controlled and minimal data exposure.

6.3 Example Use Case

Query: A patient asks the chatbot, *"What is my most recent HbA1c result?"*

- HL7 v2 system contains lab results as **ORU^R01 messages**.
- The transformation layer maps HbA1c results into a **FHIR Observation resource** with standardized codes (LOINC: 4548-4).
- API gateway fetches only the **latest lab result** and returns it in JSON.
- LLM generates a natural language response: *"Your most recent HbA1c result is 7.1%, measured on June 12, 2025."*

This ensures the LLM delivers **precise and privacy-preserving responses** without exposing unnecessary historical or unrelated records.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 8.699

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 1, January 2026

|DOI: 10.15680/IJIRSET.2026.1501004|

6.4 Conceptual Diagram

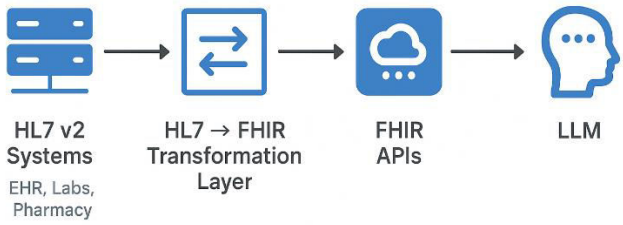

6.5 Benefits of HL7–FHIR Interoperability for GenAI

- **Selective Disclosure:** FHIR supports fine-grained queries, aligning with data minimization principles.
- **Standardized Semantics:** LOINC and SNOMED CT ensure consistent interpretation across systems.
- **Future-Proofing:** Cloud-native FHIR APIs integrate easily with AI pipelines and monitoring tools.
- **Enhanced Security:** The transformation layer acts as a **buffer zone**, reducing the risk of exposing raw HL7 data to LLMs.

This interoperability layer enables **trusted, secure, and standardized AI access** to healthcare records, ensuring compliance while unlocking the full potential of GenAI.

VII. DEPLOYMENT STRATEGY AND API SECURITY CONTROLS

Integrating LLMs with healthcare APIs requires more than interoperability—it demands a carefully designed **deployment strategy** that embeds security at every layer of the architecture. This ensures that GenAI solutions remain **scalable, compliant, and resilient** against evolving cyber threats.

7.1 Secure API Gateway Architecture

The **API gateway** acts as the control point between LLMs and healthcare systems. It enforces:

- **Authentication & Authorization:** Using OAuth 2.0 / OIDC to bind patient or clinician identity to every query.
- **Rate Limiting:** Prevents bulk extraction of PHI through iterative queries.
- **Payload Inspection:** Filters and rejects malicious or anomalous input before reaching the backend.
- **Context-Aware Access:** Dynamically adjusts permissions based on clinical role and query intent.

7.2 Zero Trust in Healthcare AI

- **"Never trust, always verify"** model applied to LLM requests.
- Each API call must be validated against **role, attribute, and purpose-based access controls**.
- Encryption-in-transit (TLS 1.3) and encryption-at-rest (AES-256) to safeguard PHI.
- Micro-segmentation isolates EHR, lab, pharmacy, and AI services, minimizing blast radius.

7.3 Monitoring and Observability

- **Audit Logging:** Every query, transformation, and LLM response is logged for compliance review.
- **Anomaly Detection:** ML models monitor for unusual query patterns (e.g., a patient portal suddenly requesting hundreds of lab results).
- **Explainability Reports:** Track how LLM responses are derived from FHIR data to ensure transparency.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 8.699

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 1, January 2026

|DOI: 10.15680/IJIRSET.2026.1501004|

7.4 Deployment Models

- **On-Premises:** Preferred for hospitals with strict PHI residency requirements.
- **Hybrid:** HL7 → FHIR transformation on-prem, with LLM access via secure cloud APIs.
- **Cloud-Native:** Fully cloud-deployed with HIPAA/GDPR-compliant hosting (Azure Health Data Services, Google Cloud Healthcare API).

VIII. PROBLEM STATEMENT AND PROPOSED SOLUTION

8.1 Problem Statement

The integration of generative AI (GenAI) and large language models (LLMs) with healthcare APIs presents a dual-edged challenge. On one side, it promises transformative potential — enabling patient-centered interactions, intelligent diagnostics, and automated administrative workflows. On the other, it exposes healthcare systems to unprecedented privacy, security, and interoperability risks.

Despite the advances in FHIR-based interoperability and zero-trust security principles, three persistent issues remain:

1. **Uncontrolled Data Exposure:**

LLMs, when interfaced with EHRs through APIs, can unintentionally access or infer sensitive patient information beyond the intended scope. This violates regulatory principles of *minimum necessary access* and creates compliance vulnerabilities under HIPAA, GDPR, and local privacy acts.

2. **Vulnerability to Prompt Injection and Adversarial Attacks:**

Traditional security models are not designed for the unstructured, context-driven nature of natural language queries. Attackers can manipulate prompts to override guardrails, extract unauthorized data, or alter the intended functionality of clinical AI assistants.

3. **Fragmented Interoperability Across Legacy Systems:**

The coexistence of HL7 v2-based legacy systems and modern FHIR architectures causes semantic inconsistencies and integration bottlenecks. Without a standardized transformation layer, GenAI systems cannot safely or accurately query clinical data across multiple platforms.

These issues collectively create an environment where patient privacy, system reliability, and regulatory compliance are at constant risk — undermining the trust required for GenAI adoption in healthcare.

8.2 Proposed Solution

To address these challenges, a **comprehensive security and interoperability framework** is proposed, integrating **zero-trust design**, **policy-driven data governance**, and **FHIR-based normalization**. The approach focuses on embedding security into every interaction between LLMs and healthcare APIs, ensuring that no single layer is solely responsible for protection.

A. Architectural Reinforcement

- **API Gateway Mediation:**

All LLM requests are routed through a secure API gateway that authenticates, authorizes, and filters data access. The gateway enforces *intent whitelisting*, *role-based filtering*, and *query validation* before forwarding any request to backend systems.

- **Prompt Shielding and Context Isolation:**

Input prompts are pre-processed using NLP-based sanitization models to detect prompt injection or adversarial content. Sensitive contexts (e.g., PHI datasets) are logically isolated from general knowledge bases to prevent data leakage or hallucination.

B. Data Governance and Compliance Controls

- **Dynamic Consent and Role-Aware Access:**

Integration of dynamic patient consent management and ABAC (Attribute-Based Access Control) ensures real-time authorization decisions based on user context, location, and purpose.

- **De-identification Pipelines:**

Sensitive identifiers are masked, tokenized, or pseudonymized at the API response layer. This ensures that only essential attributes are exposed to the LLM for reasoning, reducing re-identification risks.

- **Auditability and Explainability:**

Every LLM-mediated API call is logged, traced, and linked to a specific purpose or user. An explainability dashboard enables compliance officers to reconstruct the reasoning path behind each AI-generated recommendation.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 8.699

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 1, January 2026

|DOI: 10.15680/IJIRSET.2026.1501004|

C. Interoperability and Standardization

- **HL7–FHIR Transformation Layer:**

Legacy HL7 data is translated into FHIR resources using standardized terminologies such as LOINC and SNOMED CT. This normalization ensures semantic consistency and selective data exposure suitable for GenAI interactions.

- **Context-Aware Query Mapping:**

The transformation layer interprets natural language queries (e.g., "Show my last cholesterol test") and maps them to FHIR-compliant resources, returning structured JSON responses that align with minimal disclosure principles.

D. Continuous Threat Monitoring and Adaptation

- **Anomaly Detection Models:**

Machine learning-based observability tools analyze query patterns to detect potential data scraping or privilege escalation.

- **Adaptive Guardrails:**

Security policies dynamically adjust based on usage trends, patient risk scores, or contextual triggers (e.g., multiple failed authentications).

8.3 Outcome and Benefits

Implementing this solution delivers measurable improvements in security, compliance, and system resilience:

Challenge	Proposed Countermeasure	Expected Outcome
PHI overexposure via APIs	API-level data minimization and masking	Reduced risk of privacy violations
Prompt injection attacks	NLP-based intent validation and isolation	Prevention of unauthorized data extraction
Legacy interoperability	HL7–FHIR normalization layer	Unified, secure data exchange for AI access
Compliance traceability	Continuous audit logging and explainability	Regulatory readiness and accountability
Dynamic governance	Contextual consent and ABAC enforcement	Improved patient control and trust

This framework bridges the critical gap between **AI innovation** and **healthcare regulation**, enabling responsible, secure, and scalable GenAI adoption in clinical environments.

IX. CONCLUSION

The integration of generative AI and large language models into healthcare systems offers unprecedented opportunities to enhance patient engagement, streamline clinical workflows, and improve decision-making. However, these benefits come with significant risks, particularly concerning **privacy, security, and interoperability**.

This paper proposed a **secure, multi-layered architecture** for enabling LLM access to sensitive healthcare records through standardized APIs. The framework emphasized:

- **Controlled Data Exposure:** Ensuring that LLMs only receive the minimum necessary information.
- **Prompt Injection Mitigation:** Defending against adversarial queries through gateway-level filtering, model guardrails, and context isolation.
- **Interoperability Enablement:** Transforming legacy HL7 data into FHIR resources for standardized, AI-ready access.
- **Governance and Compliance:** Embedding HIPAA/GDPR-aligned privacy controls and auditability into every transaction.
- **Deployment Roadmap:** Implementing zero-trust principles, monitoring, and phased rollout strategies to scale securely.

As healthcare organizations adopt GenAI solutions, **responsible deployment strategies** will be critical to ensuring trust. Future research should explore **federated learning**, **privacy-preserving AI techniques**, and **formal verification methods for AI prompts**, extending beyond technical security toward a **governance-first model of AI in healthcare**.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 8.699

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 1, January 2026

|DOI: 10.15680/IJIRSET.2026.1501004|

Ultimately, the **success of GenAI in healthcare** depends not on how powerful the models are, but on **how securely and ethically they are integrated into existing clinical ecosystems**.

REFERENCES

1. V. K. Adari, "Interoperability and Data Modernization: Building a Connected Banking Ecosystem," *Int. J. Comput. Eng. Technol.*, vol. 15, no. 6, pp. 653–662, 2024
2. IBM Security. (2024). *Zero Trust in Healthcare: Implementing Data-Centric Security for AI*. White Paper.
3. Google Cloud Healthcare API Documentation. https://cloud.google.com/healthcare
4. Microsoft Azure Health Data Services. https://learn.microsoft.com/azure/healthcare-apis
5. OpenAI. (2023). *Mitigating Prompt Injection Attacks*. Technical Report.
6. NIST (2023). *AI Risk Management Framework (AI RMF 1.0)*. National Institute of Standards and Technology.
7. Health Level Seven International (HL7). *FHIR Overview*. https://hl7.org/fhir/

INTERNATIONAL STANDARD SERIAL NUMBER INDIA

INNO SPACE SJIF Scientific Journal Impact Factor

doi cross ref

निस्केयर NISCAIR

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

9940 572 462 6381 907 438 ijirset@gmail.com

www.ijirset.com

Scan to save the contact details