



## FHIR-POWERED INTELLIGENCE: CLOUD AND AI FOR NEXT-GEN HEALTHCARE

Somnath Narayan Mahale

Independent Data & AI Consultant, Pune, 411015.

\*Corresponding author: [somnath.mahale@gmail.com](mailto:somnath.mahale@gmail.com)

Received on: 30/05/2025 Accepted on: 03/08/2025 Published on: 03/10/2025

### ABSTRACT

Healthcare data has long been locked in silos fragmented across electronic health records, registries, laboratories, and insurance systems. This fragmentation slows patient care, complicates compliance, and prevents the timely use of analytics. In the U.S., the adoption of FHIR (Fast Healthcare Interoperability Resources) combined with cloud-native data architectures has begun to transform this reality, reducing onboarding time for new data sources, enabling real-time dashboards, and powering AI-ready datasets. Drawing from my experience leading end-to-end healthcare data platforms, this paper presents an architectural approach to unifying disparate healthcare datasets into a FHIR-native clinical data repository deployed on Google Cloud Platform. I illustrate implementation through concrete use cases preventive screening, immunization compliance, and AI enablement and share results that demonstrate reduced latency, compliance-by-design, and improved clinical impact. The analysis extends beyond the U.S. by examining India's digital health journey, where national initiatives like the Ayushman Bharat Digital Mission (ABDM) and National Health Claim Exchange (NHCE) are embedding FHIR at the core of a digital public good. By comparing U.S. and Indian contexts, I highlight both challenges and opportunities for scaling interoperability, compliance, and intelligence globally. The paper concludes by looking ahead to future directions—multimodal integration, federated learning, and cross-border health data exchange—positioning FHIR as the foundation for a new era of healthcare intelligence

**KEYWORDS:** FHIR, Cloud-Native Data Architectures, Interoperable Health Records, Clinical Data Repository, Real-World Evidence, Digital Health

### ABOUT THE AUTHOR

Somnath is a Data and AI strategist with more than 18 years of global experience driving digital transformation and business growth. He specializes in data engineering, Artificial Intelligence, and cloud modernization helping organizations harness technology to create measurable impact and competitive advantage. Beyond executive consulting, Somnath contributes actively to the Data & AI ecosystem through published research, peer reviews, and patents. He is also a recognized voice in thought leadership—regularly writing on articles on media. Leading professional development workshops—empowering educators, professionals, and the next generation of engineers to thrive in a data-driven future.



### INTRODUCTION

Healthcare is simultaneously one of the most data-rich and data-fragmented industries in the world. In the U.S., a single patient's information may be scattered across dozens of registries, electronic health record (EHR) systems, lab networks, and payer databases. The lack of interoperability makes it difficult to assemble a holistic view of patient care. For providers, this means delayed decisions. For regulators, inconsistent reporting. For researchers, limited access to real-world evidence.(1–3)

The result is a widening gap between data availability and data usability. While the healthcare sector produces vast amounts of digital data, much of it is “dark data”: technically present, but inaccessible for insight.(2,4)

The solution is not merely technical but architectural. FHIR, developed by HL7 International, offers a modern, API-friendly standard to represent healthcare concepts such as patients, encounters, medications, and lab results.(5,6) On its own, however, a standard is not enough. It must be combined with cloud-native architectures pipelines, distributed storage, and scalable compute—that can normalize, validate, and analyze data at scale, while ensuring governance and compliance.(7,8)

This paper demonstrates such an architecture in practice, drawing on U.S. implementations that achieved measurable impact: onboarding new registries in days instead of weeks, reducing reporting latency to seconds, and enabling machine learning for risk prediction. The discussion then extends to India's digital health mission, where national programs are adopting FHIR at scale, creating opportunities to leapfrog traditional challenges and build interoperability into the foundation of healthcare delivery.(5,9)The goal is to show how FHIR and cloud-native platforms together accelerate the journey from data silos to actionable intelligence, making healthcare systems more compliant,

more interoperable, and more patient-centered.

### BACKGROUND & RELATED WORK

#### Data Standards in Healthcare

Healthcare interoperability efforts have evolved over decades. HL7 v2 enabled message-based exchanges, while C-CDA (Consolidated Clinical Document Architecture) standardized clinical documents. Yet both approaches struggled with flexibility and adoption at scale. FHIR (Fast Healthcare Interoperability Resources) emerged as a next-generation standard, combining web technologies (REST APIs, JSON/XML) with healthcare semantics, making data exchange both developer-friendly and clinically meaningful.(5,6,10)

#### Cloud-Native Platforms in Healthcare

Cloud providers such as Google Cloud, AWS, and Azure have become critical for healthcare modernization. They offer managed services for big data processing (e.g., Dataproc, Dataflow), scalable analytics (BigQuery), and high-performance storage (Bigtable).(7,11) Prior research has demonstrated that cloud-based data lakes improve elasticity and cost efficiency, but challenges remain in governance, compliance, and semantic interoperability.(11)

#### Gaps in Literature and Practice

Most academic work has focused on standards (FHIR specifications, implementation guides) or platforms (data lakes, ETL pipelines) in isolation. Fewer studies describe real-world architectures combining both in production at scale. Moreover, literature from India is sparse—despite the country's bold initiatives like ABDM/NHCE. This paper seeks to fill that gap by combining

lived U.S. experience with India-context insights.

## DATA ARCHITECTURE

At the heart of the implementation is a FHIR-native Clinical Data Repository (CDR) deployed on Google Cloud. The end-to-end pipeline is as follows:

**Data Ingestion:** Clinical data from registries, EHRs, and labs arrive via SFTP, RPC, or APIs.

**Landing Zone (GCS):** Raw files are ingested into Google Cloud Storage (GCS).

**Transformation (Dataproc, Dataflow):** Spark and streaming pipelines parse data, map it to FHIR R4 resources, and validate structure.(7,11)

**FHIR CDR:** Canonical JSONs are stored in GCS buckets; struc-

tured equivalents are written into BigQuery (for analytics) and Bigtable (for real-time dashboards).(7)

**APIs (GKE):** RESTful FHIR APIs exposed via Kubernetes micro-services enable external systems to consume standardized bundles.

## SECURITY AND GOVERNANCE

HIPAA compliance guided the design, with encryption-at-rest, access controls, audit logs, and PHI minimization. In India, ABDM similarly enforces consent-driven data sharing, ensuring that privacy is preserved as data flows across hospitals, payors, and patients.(12)

## FHIR IMPLEMENTATION

### Mapping and Normalization

Each data source is mapped into FHIR R4 resources:




Figure 1: FHIR-native data architecture on Google Cloud Platform

- Demographics → Patient
- Encounters → Encounter
- Lab results → Observation
- Immunization → Immunization
- Medications → MedicationRequest/Medication Administration

## Validation

Validation formed a critical stage of the pipeline, ensuring that only standards-compliant FHIR resources entered the CDR. Automated routines were built using **HAPI FHIR** libraries, which provide schema-aware validation against FHIR R4 and custom profiles (e.g., U.S. Core, ABDM extensions). (5,6)

Incoming records—whether HL7, C-CDA, or raw CSV—were parsed and checked for required fields, datatype consistency, reference integrity, and controlled terminologies (SNOMED, LOINC, RxNorm). Any invalid record was quarantined into a separate GCS bucket, with error logs for remediation. This also allows data stewards and integration teams to quickly diagnose issues without polluting the downstream pipeline.

The quarantine workflow effectively acted as a firewall for bad data, preventing schema drift or malformed payloads from reaching the FHIR CDR. More importantly, it created an auditable trail of rejected records, which proved invaluable during compliance audits (HIPAA in the U.S. and emerging ABDM compliance frameworks in India).

Validated records moved seamlessly into the canonical FHIR store, giving downstream systems a high degree of confidence in data integrity. This not only reduced the cost of rework and troubleshooting but also ensured that analytics, regulatory reporting, and AI models were powered by clean, trust-

worthy inputs. By embedding validation early in the ingestion layer, the architecture enforced a “compliance-first design”, where interoperability and data quality became default outcomes rather than afterthoughts.

## Storage and APIs

Storage and API design played a dual role—balancing auditability with performance. All incoming, validated FHIR resources were preserved in their canonical JSON form within Google Cloud Storage (GCS). This immutable store functioned as a single source of truth, supporting regulatory audits and forensic traceability.

At the same time, structured versions of these FHIR resources were flattened into BigQuery tables, enabling high-speed SQL queries across billions of records for analytics, quality reporting, and AI workloads. Low-latency patient-facing applications were supported by Bigtable, optimized for sub-second scorecards and dashboards.(7)(11)

For interoperability, FHIR REST APIs were exposed via GKE-hosted micro-services (using HAPI FHIR or custom implementations). These APIs allowed downstream systems—ranging from regulatory reporting platforms and payors to ABDM/NHCX exchanges in India—to directly retrieve FHIR bundles without custom ETL pipelines. This design ensured that the same data served multiple audiences: regulators, clinicians, researchers, and insurers, all from a consistent, standards-aligned interface.


## USE CASES

### Preventive Screening

Preventive care is one of the clearest demonstrations of the value of standardized data. In the U.S., we implemented osteoporosis screening for women aged

65–85 (MIPS #39) by filtering patients by demographics, confirming visits through encounter data, and checking FHIR Observation resources for DEXA scans. The same workflow applies in India with a different focus: anemia screening among women of reproductive age. Here, the logic is identical—Patient resources define the eligible demographic, Encounter confirms the visit (at a PHC or health camp), and Observation validates the presence of an Hb test. This consistency illustrates how FHIR enables preventive care logic to be reused across geographies, reducing rework while improving public health out-

**FHIR Resource Relationships (Patient-Centric View)**



**Figure 2: FHIR-Resource Relationship (Patient-Centric View)**

comes.

Patient filter → eligible demographic.

Encounter → confirm visit (PHC or health camp).

Observation → check for Hb test results.

## IMMUNIZATION COMPLIANCE

Immunization programs are critical both for public health and for reimbursement-driven compliance. In the U.S., we tracked seasonal influenza vaccinations by consolidating records from EHRs, immunization registries, and claims. In India, similar workflows can power the Universal Immunization Programme (UIP) and COVID-19 vaccination initiatives(12). The FHIR representation is clear: Patient identifies eligibility, Immunization records vaccine details (CVX codes, administration dates), and when registry feeds are incomplete, claims data is normalized into FHIR Immunization resources. By standardizing these sources, care teams can track coverage, identify gaps, and conduct targeted outreach. The result is improved compliance, reduced duplicate reporting, and stronger population health intelligence.

Patient → eligibility.

Immunization → vaccine administration records.

Claims → normalized into Immunization when registry feeds are incomplete.

## AI/ML ENABLEMENT

Once healthcare data is standardized into FHIR and stored in analytics-ready platforms, it becomes fertile ground for AI and machine learning. In our implementations, we built curated datasets that predicted:

- Osteoporosis screening gaps, enabling outreach for preventive care
- 30-day readmission risk, improving care coordination.
- Likelihood of vaccination compliance, allowing providers to target interventions for at-risk populations.

These AI-driven insights go beyond compliance—they directly improve patient outcomes by bridging clinical workflows with predictive analytics. In India, as ABDM data lakes mature, the same approach can support AI-driven anemia detection, maternal health risk prediction, and immunization forecast-

ing, accelerating the shift from reactive to proactive care.

The architecture was evaluated not only on technical performance but also on regulatory and clinical outcomes. Registry onboarding, which previously required manual mapping and custom ETL development, was accelerated dramatically—from several weeks per source down to just a few days—thanks to the configurable ingestion and normalization framework. This speed-up meant new registries and EHR feeds could be integrated into the FHIR CDR far more rapidly, ensuring data availability kept pace with clinical and business needs. For real-time insights, the system consistently achieved dashboard refresh latencies in a few seconds by leveraging Bigtable for low-latency scorecards. This near-instant responsiveness proved valuable in operational settings where clinicians and administrators required timely updates on performance metrics. On the compliance side, the solution successfully passed a HIPAA audit with zero critical findings. HIPAA audit passed with zero critical findings,(13) demonstrating that security and governance controls (IAM, encryption, audit logging) were effective by design rather than retrofitted after implementation. This compliance-first approach reduced risk and built trust among stakeholders.

Most importantly, there was clear clinical and business impact. Standardized FHIR data enabled accurate measurement of preventive screenings and immunizations, leading to higher compliance rates and improved reimbursement outcomes under value-based care programs. In the Indian context, similar benefits would directly translate to national initiatives like ABDM and NHCX, where better data integrity can improve public health reporting, streamline claims, and support population health management.

## DISCUSSION

Several important lessons emerged from this implementation. First, interoperability is as much about governance as it is about technology.(1)(2) It is relatively straightforward to map HL7 messages or C-CDA documents into FHIR resources, but aligning hospitals, registries, insurers, and regulators on consistent policies, consent frameworks, and coding standards requires sustained governance. Without this alignment, even the most elegant pipelines cannot achieve their full potential.

Second, FHIR proved to be a powerful abstraction layer. Once data from disparate systems was normalized into FHIR, new downstream use cases multiplied—regulatory reporting, real-time dashboards, AI/ML pipelines, and even cross-system APIs could all draw from the same standardized foundation. This separation of ingestion complexity from analytic consumption simplified development and improved reusability.(6)

The implications for Pharmatech are significant. Standardized patient journeys expressed in FHIR enable stronger real-world evidence generation, which is increasingly required by regulators and payors. Similarly, the availability of longitudinal data supports faster, more precise clinical trial recruitment, reducing costs and improving representativeness. FHIR APIs also facilitate therapeutic monitoring and precision medicine, allowing pharma and life sciences companies to track treatment responses in near real time.(13)

From a global perspective, India's ABDM and NHCX initiatives illustrate how these principles can scale to the level of a digital public good.(9) By embedding FHIR at the national layer, India has the opportunity to leapfrog the incremental, often fragmented adoption path taken in the U.S., and instead build an interoperable health data ecosystem that supports care delivery, research, and innovation at population scale.

## Highlights:

- Interoperability requires governance alignment, not just technical standards.



- FHIR simplifies reuse by decoupling ingestion from analytics.
- Pharmatech benefits: real-world evidence, faster trials, precision therapies.
- India's ABDM/NHCX shows national-scale FHIR adoption potential.

## FUTURE DIRECTIONS

Looking ahead, healthcare data platforms are poised to extend FHIR into entirely new domains. A key priority will be multimodal integration, where FHIR serves as the unifying schema for not just EHR and claims data, but also genomic sequences, medical imaging, and wearable device signals.<sup>(14)</sup> This fusion of structured and unstructured data will unlock a more complete digital phenotype of patients.

Another frontier is the intersection of federated learning and generative AI. With FHIR providing standardized schemas, it becomes possible to train predictive and generative models across distributed datasets without moving sensitive patient data, thus safeguarding PHI <sup>(15)</sup> while enabling collaborative innovation. Such approaches could advance personalized medicine while preserving privacy.

Finally, cross-border interoperability is emerging as both a challenge and an opportunity. With globalization of clinical trials and increased patient mobility, <sup>(16)</sup> the ability to share FHIR-native health records across geographies could harmonize regulatory compliance, speed up drug development, and support continuity of care for patients moving between countries. The standards are in place—the next phase is governance alignment and operational execution on a global scale.

## Highlights:

- Extend FHIR to multimodal health data.
- Use federated learning and GenAI with standardized schemas.
- Enable cross-border interoperability for global trials and patient mobility.

## CONCLUSION

FHIR, when combined with cloud-native data architectures, is no longer just a compliance mechanism—it becomes the engine that transforms fragmented healthcare data into a foundation for intelligence. The U.S. experience illustrates how thoughtfully applied standards can cut onboarding time, improve latency, and enable advanced analytics. India's ABDM and NHCX take this further, proving that FHIR can operate as the backbone of a national digital health ecosystem, creating public goods at population scale.

What is emerging now is an AI-powered layer on top of FHIR. By delivering clean, interoperable, and trusted data, FHIR-native platforms make it possible to train predictive AI models for risk stratification, and even Generative AI models capable of summarizing longitudinal health journeys, drafting clinical documentation, or simulating trial cohorts. These capabilities show that once data flows freely and securely, innovation compounds—quality reporting, pharmacovigilance, precision medicine, and global clinical trials all accelerate.

Together, the U.S. journey of incremental adoption and India's leapfrogging at national scale highlight a future where healthcare is compliant, interoperable, and AI-ready by design.<sup>(9)</sup> With FHIR as the connective tissue and cloud as the enabler, the next decade will see healthcare data platforms evolve from pipelines into intelligent ecosystems, powering breakthroughs in care delivery, drug discovery, and global collaboration. Finally authors are of opinion that "FHIR + Cloud + AI is not just plumbing for healthcare data—it's the foundation for a global, intelligent health ecosystem."

## REFERENCES

1. Torab-Miandoab A, Samad-Soltani T, Jodati A, Rezaei-Hachesu P. In-

- teroperability of heterogeneous health information systems: a systematic literature review. *BMC Med Inform Decis Mak.* 2023;23(1):18. <https://doi.org/10.1186/s12911-023-02115-5>
2. Ayaz M, Pasha MF, Alzahrani MY, Budiarto R, Stiawan D. The Fast Health Interoperability Resources (FHIR) standard: systematic literature review of implementations, applications, challenges and opportunities. *JMIR Med informatics.* 2021;9(7):e21929. <https://doi.org/10.2196/21929>
3. Adler-Milstein J, Jha AK. HITECH Act drove large gains in hospital electronic health record adoption. *Health Aff.* 2017;36(8):1416–22. <https://doi.org/10.1377/hlthaff.2016.1651>
4. Raghupathi W, Raghupathi V. Big data analytics in healthcare: promise and potential. *Heal Inf Sci Syst.* 2014;2(1):3. <https://doi.org/10.1186/2047-2501-2-3>
5. HL7 International [Internet]. 2023 [cited 2025 Sep 23]. FHIR Release 4 Specification. Available from: <https://hl7.org/fhir/> [accessed on 23 September 2025].
6. Mandel JC, Kreda DA, Mandl KD, Kohane IS, Ramoni RB. SMART on FHIR: a standards-based, interoperable apps platform for electronic health records. *J Am Med Informatics Assoc.* 2016;23(5):899–908. <https://doi.org/10.1093/jamia/ocv189>
7. Zao JKK, Wu JTS, Kanyimbo K, Delizy F, Gan TT, Kuo HI, et al. Design of a Trustworthy Cloud-Native National Digital Health Information Infrastructure for Secure Data Management and Use. *Oxford Open Digit Heal.* 2024;2:oqae043. <https://doi.org/10.1093/oodh/oqae043>
8. Hong J, Morris P, Seo J. Interconnected personal health record ecosystem using IoT cloud platform and HL7 FHIR. In: 2017 IEEE international conference on healthcare informatics (ICHI). IEEE; 2017. p. 362–7. <https://doi.org/10.1109/ICHI.2017.82>
9. Sharma RS, Rohatgi A, Jain S, Singh D. The Ayushman Bharat Digital Mission (ABDM): making of India's digital health story. *CSI Trans ICT.* 2023;11(1):3–9. <https://doi.org/10.1007/s40012-023-00375-0>
10. Benson T, Grieve G. Principles of health interoperability: SNOMED CT, HL7 and FHIR. Vol. 3. Springer; 2016. <https://doi.org/10.1007/978-3-319-30370-3>
11. Jiang M, Wu L, Lin L, Xu Q, Zhang W, Wu Z. Cloud-native-based flexible value generation mechanism of public health platform using machine learning. *Neural Comput Appl.* 2023;35(3):2103–17. <https://doi.org/10.1007/s00521-022-07221-5>
12. Zahoor I, Singh S, Behl T, Sharma N, Naved T, Subramanian V, et al. Emergence of microneedles as a potential therapeutics in diabetes mellitus. *Environ Sci Pollut Res.* 2021;1–21. <https://doi.org/10.1007/s11356-021-17346-0>
13. Sherman RE, Anderson SA, Dal Pan GJ, Gray GW, Gross T, Hunter NL, et al. Real-world evidence—what is it and what can it tell us. *N Engl J Med.* 2016;375(23):2293–7. <https://doi.org/10.1056/NEJMs1609216>
14. Chahal CAA, Alahdab F, Asatryan B, Addison D, Aung N, Chung MK, et al. Data Interoperability and Harmonization in Cardiovascular Genomic and Precision Medicine. *Circ Genomic Precis Med.* 2025:e004624. <https://doi.org/10.1161/CIRCGEN.124.004624>
15. Kaissis GA, Makowski MR, Rückert D, Braren RF. Secure, privacy-preserving and federated machine learning in medical imaging. *Nat Mach Intell.* 2020;2(6):305–11. <https://doi.org/10.1038/s42256-020-0186-1>
16. Nalin M, Baroni I, Faiella G, Romano M, Matriciano F, Gelenbe E, et al. The European cross-border health data exchange roadmap: Case study in the Italian setting. *J Biomed Inform.* 2019;94:103183. <https://doi.org/10.1016/j.jbi.2019.103183>