WCRG Global Research WORLD COMET RESEARCH GROUP

International Journal of Data Engineering and Intelligent Computing
Vol. 9, No. 3, 2026
Doi: 10.67228/30715717/IJDEIC-V9I3P104
PP. 64-102

Original Article

Design and Evaluation of a Scalable FHIR-Based Data Engineering Architecture for Integrating Heterogeneous Electronic Health Records

Venkata Manvitha Ala

Independent Researcher, United States of America (USA).

Received: 07-10-2026 Revised: 07-31-2026 Accepted: 08-11-2026 Published: 08-14-2026

ABSTRACT

The rapid growth of electronic health record (EHR) systems has created significant challenges in integrating heterogeneous healthcare data generated from diverse clinical environments. Although Fast Healthcare Interoperability Resources (FHIR) provides a standardized approach for healthcare data exchange, existing solutions often focus on resource transformation while providing limited evaluation of information preservation, scalability, observability, and reliability. This study proposes SCALE-FHIR, a scalable FHIR-based data engineering architecture for integrating heterogeneous EHR sources through modular ingestion, a Canonical Clinical Event Model (CCEM), semantic harmonization, FHIR resource generation, validation, lineage tracking, and fault-handling mechanisms. The architecture supports multiple healthcare data representations, including CSV, C-CDA, and relational EHR structures, and transforms them into standardized FHIR R4 resources. The proposed framework is evaluated using Synthea synthetic records and MIMIC-IV clinical data through experiments measuring FHIR conformance, mapping completeness, field fidelity, information loss, throughput, latency, scalability, fault detection, and recovery performance. The evaluation demonstrates that SCALE-FHIR enables reliable transformation of heterogeneous healthcare data while maintaining clinical information integrity and supporting scalable processing. The contributions of this work include an end-to-end healthcare data engineering architecture, a fidelity-aware interoperability evaluation framework, integrated observability and fault recovery mechanisms, and a reproducible benchmark methodology for assessing FHIR-based integration systems. SCALE-FHIR provides a foundation for scalable healthcare analytics, research platforms, and interoperable digital health ecosystems.

KEYWORDS

Healthcare Data Engineering, Fast Healthcare Interoperability Resources (FHIR), Electronic Health Records (EHR) Integration, Semantic Interoperability, Scalable Data Architecture, Clinical Data Harmonization.

1. INTRODUCTION

1.1. Background and Motivation

The rapid adoption of electronic health record (EHR) systems has significantly increased the volume and diversity of healthcare data generated from hospitals, laboratories, medical devices, and research platforms. These data provide valuable opportunities for clinical research, population health management, and data-driven decision support. However, healthcare information remains fragmented across independent systems that frequently use different data structures, storage models, and representation formats. This fragmentation limits effective data exchange and prevents healthcare organizations from fully utilizing available clinical information. Although interoperability standards such as Fast Healthcare Interoperability Resources (FHIR) have improved standardized healthcare data exchange, successful integration requires more than converting data into FHIR resources. Healthcare organizations require scalable data engineering architectures capable of integrating heterogeneous sources, preserving clinical meaning, validating transformations, and supporting reliable processing.

Therefore, there is a need for comprehensive FHIR-based architectures that address not only interoperability but also scalability, data quality, transparency, and operational reliability.

1.2. Challenges of Heterogeneous EHR Integration

Integrating heterogeneous EHR systems remains challenging due to differences in data structures, formats, and semantic representations. Healthcare systems commonly store clinical information using different database schemas, making direct integration difficult without complex transformation processes.

Healthcare data are also distributed across multiple formats, including relational databases, CSV files, clinical documents such as C-CDA, and standardized FHIR resources. Supporting these diverse representations requires flexible data pipelines capable of adapting to multiple source systems without creating isolated transformation workflows.

In addition, semantic inconsistencies create challenges when different healthcare organizations use different terminology systems, coding standards, and representations for similar clinical concepts. Without effective semantic harmonization, transformation processes may introduce information loss or reduce the clinical usefulness of integrated data. These challenges demonstrate that healthcare interoperability requires advanced data engineering approaches that combine structural transformation, semantic mapping, and quality preservation.

1.3. Limitations of Existing Approaches

Recent studies have demonstrated the effectiveness of FHIR-based approaches for healthcare data transformation and interoperability. Existing solutions have successfully addressed clinical document conversion, research dataset integration, and domain-specific data standardization. However, several limitations remain.

First, many approaches primarily evaluate structural FHIR compliance without measuring whether clinically important information is preserved during transformation. A resource may satisfy FHIR requirements while still losing meaningful source information.

Second, scalability evaluation remains limited. Many existing pipelines are assessed using specific datasets or clinical scenarios without comprehensive benchmarking under increasing workloads, streaming conditions, or different processing configurations.

Third, observability is often insufficient. Without detailed lineage tracking, monitoring, and error classification, identifying transformation failures and maintaining reliable healthcare data pipelines becomes difficult.

Finally, fault tolerance and recovery mechanisms are frequently overlooked, despite being essential for operational healthcare systems that must process continuous clinical data reliably.

1.4. Research Gap

Although existing research has advanced FHIR adoption and healthcare interoperability, current solutions generally address individual challenges such as transformation, semantic mapping, or domain-specific integration. A unified architecture that simultaneously evaluates heterogeneous data integration, FHIR conformance, information fidelity, and scalability, observability, and resilience remains insufficiently explored.

This study addresses this gap by proposing and evaluating an end-to-end FHIR-based data engineering architecture designed for scalable, reliable, and transparent healthcare data integration.

1.5. Aim and Objectives

The aim of this study is to design and evaluate a scalable FHIR-based data engineering architecture for integrating heterogeneous electronic health records.

The objectives are to:

- Develop a modular architecture for integrating heterogeneous EHR sources using FHIR-based transformation workflows.
- Establish a source-independent transformation approach using a Canonical Clinical Event Model and semantic mapping mechanisms.
- Evaluate information preservation during source-to-FHIR transformation.
- Benchmark scalability under large-scale batch and streaming healthcare workloads.

ISSN: 3071-5717
PP: 64-102.

- Assess observability, fault detection, and recovery capabilities under controlled failure scenarios.

1.6. Research Questions

This study investigates the following research questions:

- RQ1: How effectively can a scalable FHIR-based architecture transform heterogeneous EHR sources into valid FHIR R4 resources?
- RQ2: To what extent does the architecture preserve clinically meaningful information during transformation?
- RQ3: How scalable is the architecture under increasing healthcare data volumes and streaming workloads?
- RQ4: How effectively can lineage and monitoring mechanisms identify transformation errors?
- RQ5: How resilient is the architecture under data and infrastructure failures?
- RQ6: Can the proposed evaluation framework support reproducible assessment of FHIR-based data engineering architectures?

1.7. Contributions

This study makes five contributions:

- Architecture Contribution: A scalable FHIR-based architecture (SCALE-FHIR) is proposed for integrating heterogeneous EHR sources through modular ingestion, normalization, mapping, validation, and storage components.
- Fidelity-Aware Evaluation Contribution: An evaluation framework is introduced that measures mapping completeness, field fidelity, and information loss beyond basic FHIR validation.
- Observability Contribution: The architecture incorporates lineage tracking, monitoring, and fault classification to improve transparency and reliability.
- Scalability Benchmark Contribution: A benchmark framework evaluates batch and streaming performance under increasing workloads and processing configurations.
- Reproducible Evaluation Contribution: The study defines datasets, metrics, acceptance criteria, and fault-injection procedures for repeatable evaluation.

1.8. Paper Organization

The remainder of this paper is organized as follows. Section 2 reviews existing research on healthcare data engineering, FHIR interoperability, semantic harmonization, and related integration pipelines. Section 3 presents the proposed SCALE-FHIR architecture. Section 4 describes the datasets, experimental design, evaluation metrics, and benchmark procedures. Section 5 presents the results, followed by discussion in Section 6. Section 7 outlines limitations and future work, and Section 8 concludes the paper.

ISSN: 3071-5717
PP: 64-102.

2. BACKGROUND AND RELATED WORK

2.1. Healthcare Data Engineering and EHR Heterogeneity

The increasing adoption of electronic health record (EHR) systems has generated large volumes of clinical data distributed across heterogeneous healthcare environments. However, healthcare organizations commonly operate multiple information systems with differences in data structures, terminology usage, storage models, and exchange mechanisms. These variations create significant challenges for integrating clinical information across institutions and developing scalable secondary-use applications.

Torab-Miandoab et al. identified interoperability between heterogeneous health information systems as a persistent challenge caused by differences in architecture, semantics, and data representation. Their review emphasized that successful healthcare data integration requires not only technical connectivity but also mechanisms capable of maintaining meaningful clinical interpretation across different systems.

Similarly, Martens et al. demonstrated that modern healthcare environments contain diverse data sources, including electronic medical records, medical devices, and operational systems, requiring flexible integration approaches capable of handling continuously generated clinical information. The increasing use of smart hospital technologies further increases the need for data engineering architectures that support structured collection, transformation, and exchange of healthcare information.

Ambalavanan et al. further highlighted that building healthcare data ecosystems requires addressing multiple challenges, including inconsistent data structures, incomplete information, limited standard adoption, and difficulties achieving sustainable interoperability. These findings indicate that healthcare interoperability is not solely a data format problem but a broader data engineering challenge involving ingestion, transformation, validation, quality management, and reliable access.

2.2. HL7 FHIR and Healthcare Interoperability

Fast Healthcare Interoperability Resources (FHIR), developed by HL7 International, has become one of the most widely adopted standards for healthcare data exchange. FHIR provides a resource-based approach in which clinical concepts are represented through modular resources with defined structures, relationships, and exchange mechanisms. FHIR Release 4 (R4) introduced a stable foundation widely adopted in healthcare applications, while FHIR Release 5 (R5) provides additional improvements and expanded capabilities.

Although FHIR improves interoperability by providing standardized representations, successful implementation requires appropriate mapping between existing healthcare data structures

ISSN: 3071-5717
PP: 64-102.

and FHIR resource models. Griffin et al. examined real-world healthcare applications using FHIR and demonstrated that implementations vary considerably in technical architecture, clinical purpose, and deployment characteristics. Their findings indicate that FHIR adoption requires careful consideration of implementation strategies rather than simply converting existing data into FHIR formats.

Duda et al. reviewed FHIR-based tools and initiatives supporting clinical research and identified growing adoption of FHIR for research data exchange, clinical decision support, and healthcare application development. However, they also highlighted continuing challenges related to implementation complexity, resource mapping, and maintaining consistency across different healthcare environments.

Therefore, while FHIR provides a standardized interoperability framework, it does not independently resolve challenges associated with heterogeneous source systems, semantic differences, data quality variation, or large-scale processing requirements. Additional data engineering mechanisms are required to transform diverse healthcare data sources into reliable FHIR-based resources.

2.3. FHIR Mapping and Semantic Harmonization

A major challenge in FHIR adoption is transforming existing healthcare data into semantically meaningful FHIR resources. Healthcare systems frequently use different schemas, coding systems, and terminology standards, requiring mapping approaches capable of preserving clinical meaning during transformation.

Bossenko et al. investigated the transformation of HL7 Clinical Document Architecture (CDA) data into FHIR using reusable visual mapping components. Their work demonstrated that reusable mapping approaches can improve transformation efficiency and reduce manual development effort. However, mapping remains dependent on accurate interpretation of source structures and clinical concepts.

De Mello et al. conducted a systematic review of semantic interoperability approaches in health record standards and identified terminology alignment and consistent representation of clinical concepts as major barriers to effective data exchange. Their findings demonstrate that structural compatibility alone is insufficient for achieving meaningful interoperability.

Zhang et al. explored cross-standard healthcare data harmonization by focusing on semantic relationships between data elements. Their work emphasized the importance of aligning concepts across different healthcare standards rather than relying solely on direct attribute matching.

Similarly, Xiao et al. proposed an approach combining FHIR RDF and the Observational Medical Outcomes Partnership (OMOP) Common Data Model to support clinical knowledge graph

ISSN: 3071-5717
PP: 64-102.

development. This demonstrates the potential of combining healthcare standards; however, such approaches often focus on specific interoperability objectives rather than providing complete end-to-end data engineering pipelines.

These studies demonstrate significant progress in FHIR mapping and semantic harmonization. However, challenges remain in evaluating mapping completeness, measuring information preservation, and maintaining transformation transparency across heterogeneous healthcare sources.

2.4. Existing FHIR-Based Data Engineering Pipelines

Recent research has explored practical architectures for transforming and integrating healthcare data using FHIR-based pipelines. Marfoglia et al. proposed a modular FHIR-driven transformation pipeline consisting of input processing, refinement, mapping, validation, and export stages. Their work demonstrated the importance of modular architecture design for healthcare data standardization.

Carbonaro et al. extended this concept by developing a FHIR-based transformation pipeline for real-world oncology data, demonstrating the feasibility of converting large clinical datasets into research-ready representations. However, their evaluation primarily focused on transformation effectiveness within a specific clinical domain.

Barret et al. introduced an interoperability-aware healthcare metadata pipeline designed to support federated analyses. Their work addressed distributed data environments but focused mainly on interoperability-aware metadata processing rather than comprehensive evaluation of transformation fidelity and operational resilience.

Fecho et al. developed a FHIR-based pipeline for integrating geospatial and spatiotemporal clinical research data. While this demonstrated the flexibility of FHIR for specialized research applications, the approach was designed around a specific use case rather than general heterogeneous EHR integration.

Iancu et al. investigated large-scale integration of DICOM metadata into FHIR for medical research. Their work highlights the importance of extending FHIR beyond traditional clinical records, but also demonstrates the complexity of integrating additional healthcare data modalities.

Collectively, these studies establish that FHIR-based pipelines can successfully support healthcare data integration. However, existing approaches generally evaluate specific transformation scenarios and provide limited assessment of combined interoperability, fidelity, scalability, observability, and fault tolerance.

2.5. Data Quality and Information Preservation

Successful healthcare data integration requires more than generating technically valid FHIR resources. Maintaining data quality and preserving clinically relevant information throughout transformation are critical requirements.

App et al. investigated data quality and completeness in oncology real-world data and demonstrated that information available in source EHR systems may not always be preserved in downstream standardized representations. Their findings highlight the risk that technically valid transformations may still result in incomplete clinical information.

Lewis et al. reviewed electronic health record data quality assessment approaches and identified completeness, correctness, consistency, and reliability as essential dimensions for evaluating healthcare datasets. Their work emphasizes the need for systematic quality assessment before data can be used for research or analytics. Penev et al. similarly examined EHR data quality assessment methods and found that healthcare datasets require multidimensional evaluation frameworks because quality problems can emerge from different stages of data collection, storage, and transformation.

Razzaghi et al. demonstrated the importance of evaluating whether EHR data are fit for specific research purposes. Their findings reinforce that data quality assessment must consider the intended use of transformed healthcare information. These studies establish the importance of evaluating data quality during healthcare data transformation. However, existing FHIR integration approaches often emphasize structural interoperability while providing limited evaluation of source-to-target fidelity and information loss.

2.6. Comparative Analysis of Existing Approaches

Table 1 summarizes representative FHIR-based healthcare data integration approaches and highlights their primary objectives, evaluation focus, and remaining limitations.

Table 1: Comparison of Existing FHIR-Based EHR Integration Approaches

Study	Primary Objective	Data Source / Domain	FHIR Role	Main Contribution	Evaluation Focus	Remaining Limitation
Bossenko et al. (2024)	Transform HL7 CDA documents into FHIR resources using reusable mapping components	Clinical documents (CDA)	CDA-to-FHIR transformation	Introduced reusable visual mapping components to improve FHIR transformation efficiency	Mapping feasibility and interoperability	Limited evaluation of large-scale processing, data fidelity, and operational resilience

Bennett et al. (2023)	Convert MIMIC-IV clinical data into interoperable FHIR resources	MIMIC-IV critical care and hospital data	FHIR representation of EHR datasets	Developed MIMIC-IV-on-FHIR to improve accessibility and standardized data exchange	Dataset conversion and FHIR resource availability	Primarily focused on dataset conversion rather than complete pipeline scalability and fault evaluation
Xiao et al. (2022)	Integrate FHIR RDF with OMOP CDM for clinical knowledge graph development	Clinical research datasets	FHIR-based semantic representation	Demonstrated integration between healthcare standards for knowledge representation	Semantic interoperability and data modeling	Does not provide comprehensive evaluation of transformation performance or operational monitoring
Marfoglia et al. (2025)	Develop a modular FHIR-driven transformation pipeline for clinical data standardization	Real-world clinical datasets	FHIR transformation and standardization	Proposed modular architecture for converting heterogeneous clinical data into FHIR	Transformation workflow and standardization	Limited evaluation of scalability, information preservation, and failure recovery
Carbonaro et al. (2025)	Transform oncology real-world data into research-ready FHIR representations	Oncology clinical data	FHIR-based research data preparation	Demonstrated practical application of FHIR pipelines in a real clinical domain	Domain-specific transformation effectiveness	Limited generalizability beyond oncology and limited scalability benchmarking
Barret et al. (2025)	Develop interoperability-aware metadata pipelines for federated healthcare analysis	Healthcare metadata sources	FHIR-aware data integration	Supported federated analysis through standardized metadata processing	Metadata interoperability	Focuses on metadata exchange rather than full clinical data transformation lifecycle

ISSN: 3071-5717
PP: 64-102.

Fecho et al. (2025)	Integrate geospatial and spatiotemporal clinical research information	Clinical research datasets	FHIR-based data integration	Extended FHIR usage for specialized research data integration	Research data integration capability	Limited assessment of heterogeneous EHR sources and system-level reliability
Iancu et al. (2024)	Integrate medical imaging metadata into FHIR for research applications	DICOM and clinical imaging data	DICOM-to-FHIR integration	Demonstrated extension of FHIR for medical imaging interoperability	Metadata transformation and integration	Focuses on imaging data and does not address broader EHR pipeline requirements
Torab-Miandoab et al. (2023)	Review interoperability challenges across heterogeneous health information systems	Healthcare information systems	Interoperability framework analysis	Identified technical and semantic barriers affecting healthcare integration	Literature-based interoperability assessment	Does not propose or experimentally evaluate an end-to-end architecture
De Mello et al. (2022)	Analyze semantic interoperability approaches in healthcare standards	Health record standards	Semantic mapping and harmonization	Identified challenges in maintaining clinical meaning across standards	Semantic interoperability review	Does not evaluate scalable implementation strategies
Griffin et al. (2022)	Analyze real-world healthcare applications implementing FHIR	FHIR-enabled healthcare applications	FHIR implementation assessment	Characterized practical FHIR adoption patterns and implementation challenges	Clinical and technical implementation analysis	Does not provide a scalable data engineering architecture
Duda et al. (2022)	Review FHIR-based tools supporting clinical research	Clinical research initiatives	FHIR research interoperability	Summarized FHIR adoption for research workflows	Tool and initiative analysis	Limited focus on transformation fidelity and pipeline performance

ISSN: 3071-5717
PP: 64-102.

Table 1 demonstrates that previous FHIR-based integration studies have significantly advanced healthcare interoperability through resource transformation, semantic mapping, and domain-specific data pipelines. However, existing approaches typically address isolated challenges, such as CDA-to-FHIR conversion, clinical dataset standardization, or specialized research integration. Few studies evaluate the complete healthcare data engineering lifecycle, including heterogeneous source ingestion, semantic transformation, information preservation, large-scale performance, operational observability, and fault recovery within a single architecture.

The comparison demonstrates that existing studies have made substantial contributions in areas such as FHIR transformation, semantic mapping, clinical research integration, and domain-specific data pipelines. However, most approaches address individual aspects of interoperability rather than evaluating the complete lifecycle of healthcare data engineering.

2.7. Identified Research Gap

Existing research demonstrates that FHIR provides a strong foundation for healthcare interoperability and that specialized pipelines can successfully transform heterogeneous clinical data into standardized representations. However, current approaches commonly evaluate individual components, such as transformation accuracy, semantic mapping, or domain-specific integration, without jointly assessing interoperability, data fidelity, scalability, observability, and resilience.

Therefore, a unified data engineering architecture that integrates heterogeneous EHR sources while systematically evaluating structural conformance, information preservation, processing performance, fault handling, and reproducibility remains insufficiently explored. This research addresses this gap by proposing and evaluating SCALE-FHIR as an end-to-end architecture designed for scalable, reliable, and transparent FHIR-based healthcare data integration.

3. PROPOSED SCALE-FHIR ARCHITECTURE

3.1. Architecture Design Requirements

The proposed Scalable, Conformance-Aware, Lineage-Enabled FHIR Data Engineering Architecture (SCALE-FHIR) is designed to address the limitations identified in existing healthcare interoperability approaches. Although previous studies have demonstrated the feasibility of transforming heterogeneous electronic health record (EHR) data into Fast Healthcare Interoperability Resources (FHIR), many existing solutions focus on specific transformation tasks or clinical domains without jointly considering scalability, information preservation, operational visibility, and system reliability. Therefore, SCALE-FHIR was designed around six fundamental requirements: interoperability, scalability, data fidelity, observability, fault tolerance, and reproducibility.

The interoperability requirement ensures that heterogeneous healthcare data originating from different source systems, formats, and schemas can be transformed into standardized FHIR resources

ISSN: 3071-5717
PP: 64-102.

while maintaining semantic consistency. The architecture therefore separates source-specific processing from FHIR generation through intermediate normalization and mapping components.

The scalability requirement addresses the need to process increasing healthcare data volumes without significant degradation in performance. Healthcare environments generate continuously expanding clinical records, requiring architectures capable of supporting both large-scale batch processing and near-real-time event streams.

The data fidelity requirement focuses on preserving clinically meaningful information during transformation. Unlike approaches that evaluate interoperability only through successful FHIR validation, SCALE-FHIR distinguishes between structural conformance and preservation of source-level information. This enables assessment of whether transformed resources accurately represent the original clinical data.

The observability requirement ensures that each transformation stage can be monitored and analyzed. The architecture incorporates lineage tracking, performance monitoring, and error reporting mechanisms to identify where information loss or processing failures occur.

The fault tolerance requirement enables continued operation when invalid records, infrastructure failures, or temporary service interruptions occur. Rather than terminating the entire pipeline, SCALE-FHIR isolates problematic records and supports recovery mechanisms.

Finally, the reproducibility requirement ensures that transformation rules, software configurations, datasets, and evaluation procedures are explicitly controlled and versioned. This allows independent researchers to reproduce the architecture evaluation under comparable conditions.

3.2. Overall SCALE-FHIR Architecture

SCALE-FHIR follows a layered data engineering architecture that separates healthcare data ingestion, normalization, transformation, and validation, storage, and access operations. The architecture consists of seven primary processing layers: heterogeneous source systems, source adapters, data ingestion, Canonical Clinical Event Model (CCEM), semantic harmonization and FHIR mapping, validation and quality assessment, and FHIR-based storage and access.

The complete data flow begins with heterogeneous EHR sources, including structured relational databases, CSV-based exports, clinical document formats, and synthetic healthcare datasets. Source adapters convert these different representations into a unified internal structure while preserving important metadata associated with the original records.

Following ingestion, healthcare information is transformed into the CCEM, which provides a source-independent representation of clinical events. This intermediate layer separates source

ISSN: 3071-5717
PP: 64-102.

complexity from FHIR transformation logic and enables consistent processing across different healthcare systems.

The normalized clinical events then pass through semantic harmonization and mapping components, where source attributes are aligned with FHIR resources, terminology systems, and defined transformation rules. Generated FHIR resources subsequently undergo validation to verify structural conformance and evaluate information preservation.

Validated resources are stored within a FHIR repository, where they can be accessed through standardized interfaces, including REST-based APIs and bulk data exchange mechanisms.

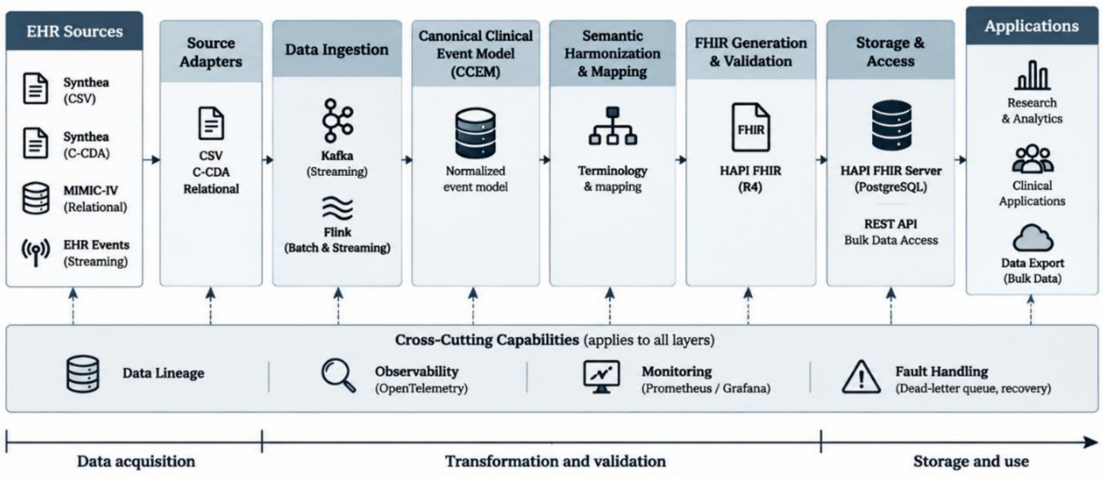

Fig 1 - Overall SCALE-FHIR Architecture Illustrates the Complete Architecture, Including the Core Processing Layers and the Cross-Cutting Capabilities of Lineage Tracking, Observability, and Fault Handling.

3.3. Heterogeneous EHR Source Layer

The heterogeneous source layer represents the diverse healthcare information systems that generate clinical data. SCALE-FHIR is designed to support multiple source representations rather than relying on a single predefined input format.

Two primary datasets are considered during evaluation: Synthea and MIMIC-IV. Synthea provides synthetic patient records generated from clinical models and supports multiple output formats, including CSV, C-CDA, and FHIR. This enables controlled evaluation of transformation consistency because the same synthetic patients can be represented through different source structures.

MIMIC-IV provides a large-scale de-identified clinical dataset derived from real hospital information systems. Unlike synthetic datasets, MIMIC-IV contains realistic variations in clinical documentation, terminology usage, and relational data organization. These characteristics make it suitable for evaluating the robustness of healthcare data engineering workflows.

The architecture supports structured CSV files, clinical document formats such as Consolidated Clinical Document Architecture (C-CDA), and relational EHR database structures. Instead of developing independent transformation pipelines for each source format, SCALE-FHIR uses dedicated source adapters that convert different inputs into a common representation before transformation.

This design improves maintainability because changes in one source system do not require redesigning the complete FHIR transformation workflow.

3.4. Data Ingestion Layer

The data ingestion layer provides the connection between healthcare data sources and downstream transformation components. SCALE-FHIR supports both batch-oriented and streaming-oriented healthcare data processing.

For event-driven healthcare data exchange, Apache Kafka is used as the messaging infrastructure. Kafka enables distributed event ingestion through partitioned topics, allowing multiple healthcare data streams to be processed concurrently. Within SCALE-FHIR, different event categories can be separated into dedicated streams, including raw clinical events, transformed resources, validation failures, and audit information.

For large-scale transformation and processing, Apache Flink is used as the primary distributed processing engine. Flink supports stateful processing, parallel execution, checkpoint-based recovery, and both batch and stream processing models. This capability allows SCALE-FHIR to evaluate healthcare data transformation under different workload conditions.

Batch ingestion supports historical EHR processing, such as converting archived clinical datasets into FHIR representations. Streaming ingestion supports continuous processing scenarios where new clinical events must be transformed and validated with minimal delay.

By supporting both processing modes, SCALE-FHIR provides a flexible foundation for healthcare organizations requiring retrospective data integration and real-time interoperability capabilities.

3.5. Canonical Clinical Event Model (CCEM)

A central component of SCALE-FHIR is the Canonical Clinical Event Model (CCEM), which acts as an intermediate representation between source systems and FHIR resources.

ISSN: 3071-5717
PP: 64-102.

Traditional transformation approaches often create direct mappings between each source format and FHIR. However, this approach becomes difficult to maintain as healthcare environments contain multiple systems with different schemas. SCALE-FHIR addresses this limitation by separating source interpretation from FHIR generation.

Each source adapter converts incoming records into CCEM objects containing standardized clinical event attributes, including patient identifiers, event types, timestamps, codes, values, units, source metadata, and transformation identifiers.

The CCEM enables source-independent transformation because multiple input formats can be processed through the same downstream mapping pipeline. For example, CSV-based laboratory information and relational laboratory tables from different EHR systems can be represented using the same clinical event structure before being transformed into FHIR Observation resources.

The model also preserves metadata required for auditing and quality evaluation. Source identifiers, schema versions, timestamps, and transformation information are maintained throughout the pipeline.

Additionally, each clinical event receives a lineage identifier, allowing individual records to be traced from the original source through transformation, validation, and storage.

3.6. Semantic Harmonization and Mapping Engine

The semantic harmonization and mapping engine converts normalized clinical events into standardized FHIR representations. This component contains a version-controlled mapping registry that defines relationships between source attributes and corresponding FHIR elements.

Each mapping rule specifies the source field, target FHIR resource, transformation logic, terminology requirements, and mapping version. Maintaining mappings as explicit configuration objects improves transparency and allows changes to be tracked throughout the research lifecycle.

Terminology normalization is performed where appropriate to improve semantic consistency. Source clinical codes may be aligned with recognized healthcare terminology systems, while preserving original source values when direct normalization is unavailable.

The mapping engine evaluates transformation outcomes by distinguishing successful mappings, partial mappings, unmapped elements, and invalid transformations. This approach allows the architecture to measure mapping completeness rather than reporting only successful resource generation.

ISSN: 3071-5717
PP: 64-102.

Version-controlled transformation rules also support reproducibility because identical source data can be transformed using the same mapping configuration.

3.7. FHIR Resource Generation and Validation Layer

The FHIR generation layer converts harmonized clinical events into standardized FHIR R4 resources. FHIR R4 is selected as the primary implementation target because of its extensive adoption across healthcare interoperability environments and availability of supporting implementation guides.

SCALE-FHIR uses HAPI FHIR libraries to generate and validate FHIR resources programmatically. This reduces errors associated with manual JSON construction and ensures compatibility with FHIR resource structures.

The evaluation focuses on seven core FHIR resources:

- Patient, representing demographic information;
- Encounter, representing healthcare visits and admissions;
- Condition, representing diagnoses and clinical problems;
- Observation, representing measurements, laboratory results, and clinical findings;
- Procedure, representing performed clinical interventions;
- MedicationRequest, representing medication orders;
- MedicationAdministration, representing administered medications.

Generated resources undergo two complementary evaluations. First, structural validation determines whether resources conform to FHIR requirements. Second, fidelity assessment compares transformed resources against source information to identify potential information loss.

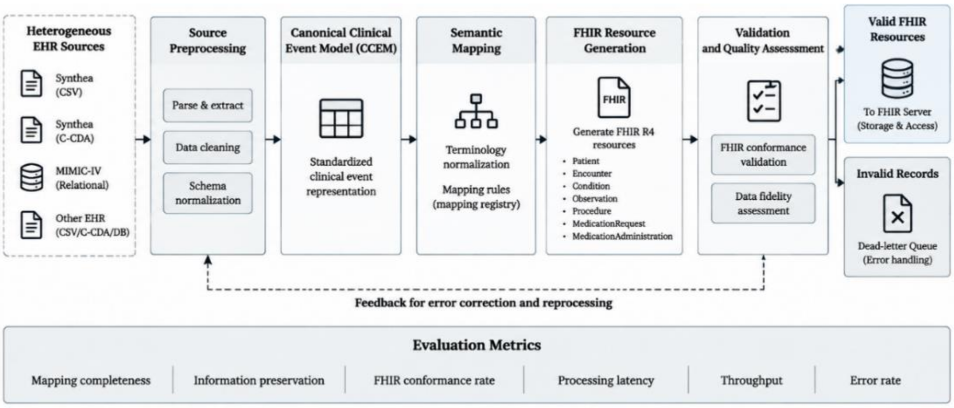

Fig 2 - Source-To-FHIR Transformation and Validation Workflow Presents the Detailed Transformation Pathway from Heterogeneous Healthcare Data Sources Through Normalization, Mapping, FHIR Generation, and Validation.

3.8. Data Lineage, Observability and Fault Handling

Observability is integrated throughout SCALE-FHIR to provide visibility into transformation behavior and system performance. Unlike traditional pipelines that evaluate only final outputs, SCALE-FHIR records processing information across each stage.

Open Telemetry is used to collect distributed traces, logs, and metrics across architecture components. These telemetry signals allow researchers to identify processing delays, transformation failures, and resource bottlenecks.

Performance monitoring is supported through Prometheus, which collects quantitative metrics including throughput, latency, resource utilization, validation failures, and pipeline errors. Grafana provides visualization dashboards for analysing operational behaviour during experiments.

Fault handling is implemented through a dead-letter queue mechanism. Records that fail validation, contain incomplete information, or encounter transformation errors are isolated rather than stopping the complete pipeline. Each failed record retains its lineage identifier, error category, processing stage, and timestamp.

Recovery mechanisms based on distributed processing checkpoints and event replay allow SCALE-FHIR to continue operation following temporary system failures.

3.9. Storage and Access Layer

The final layer provides standardized storage and access to generated FHIR resources. SCALE-FHIR uses the HAPI FHIR Server as the FHIR repository and PostgreSQL as the underlying persistence system.

The repository stores validated FHIR resources while maintaining relationships between clinical entities through standardized FHIR references. This enables downstream applications, analytics platforms, and research systems to access transformed healthcare information through standardized interfaces.

The architecture supports REST-based FHIR interactions for resource retrieval, searching, and exchange. Additionally, FHIR Bulk Data Access mechanisms are considered for large-scale export scenarios where healthcare datasets must be transferred for secondary analysis.

Through this layered design, SCALE-FHIR provides an end-to-end healthcare data engineering framework that integrates heterogeneous EHR sources while addressing interoperability, fidelity, scalability, observability, and reliability requirements.

ISSN: 3071-5717
PP: 64-102.

4. MATERIALS AND METHODS

4.1. Research Design

This study adopts a design-and-evaluation methodology to develop and systematically assess the proposed SCALE-FHIR architecture. The research approach combines software architecture design, healthcare data engineering, and controlled experimental evaluation to determine whether the proposed framework can effectively integrate heterogeneous electronic health record (EHR) sources while maintaining interoperability, data fidelity, scalability, and operational reliability.

The experimental methodology follows a quantitative evaluation approach in which the architecture is implemented using standardized healthcare data formats, open-source technologies, and reproducible configurations. The evaluation is structured around predefined research questions and hypotheses derived from the identified limitations of existing FHIR-based integration approaches.

The experiments assess five major aspects of SCALE-FHIR: (1) FHIR structural conformance, (2) preservation of source-level clinical information, (3) processing scalability under increasing workloads, (4) observability and fault detection capability, and (5) recovery performance following data and infrastructure failures. Both batch-oriented and streaming-based experiments are conducted to evaluate the architecture under different healthcare data processing scenarios.

4.2. Experimental Datasets

The evaluation uses two complementary healthcare datasets: Synthea and MIMIC-IV. Synthea provides a controlled synthetic environment where identical patient populations can be represented in multiple data formats, while MIMIC-IV provides a large-scale real-world de-identified clinical dataset for assessing the robustness of the architecture under realistic healthcare data conditions.

4.2.1. Synthea Dataset

Synthea is an open-source synthetic patient generator that produces realistic but de-identified clinical records based on computational models of disease progression and healthcare encounters. Unlike manually curated datasets, Synthea enables controlled generation of patient populations with known clinical histories, making it suitable for evaluating healthcare data transformation pipelines.

In this study, Synthea is used because it provides multiple output representations, including comma-separated values (CSV), Consolidated Clinical Document Architecture (C-CDA), and Fast Healthcare Interoperability Resources (FHIR). These different representations enable controlled assessment of how effectively SCALE-FHIR handles heterogeneous source formats originating from the same underlying clinical events.

The availability of native FHIR outputs also provides a reference representation for evaluating transformation fidelity. Source records generated as CSV and C-CDA are processed through SCALE-FHIR and compared against

The corresponding FHIR representation to assess mapping completeness, structural consistency, and information preservation. Because the patient generation process is controlled, Synthea enables repeatable experiments where transformation performance can be evaluated independently of unpredictable clinical documentation variations.

4.2.2. MIMIC-IV Dataset

MIMIC-IV is a large-scale, publicly available, de-identified clinical dataset derived from real hospital information systems. It contains extensive clinical information, including patient demographics, hospital admissions, diagnoses, procedures, laboratory measurements, medication records, and clinical events.

Unlike synthetic datasets, MIMIC-IV reflects the complexity of real healthcare environments, including variations in coding practices, incomplete documentation, heterogeneous relational structures, and missing clinical attributes. These characteristics make it suitable for evaluating the ability of SCALE-FHIR to process realistic healthcare data challenges.

This study uses structured MIMIC-IV relational data as the primary source representation. The publicly available MIMIC-IV-on-FHIR representation is used as a reference framework for evaluating resource alignment and transformation consistency.

The combination of MIMIC-IV and MIMIC-IV-on-FHIR allows evaluation of whether SCALE-FHIR can transform relational healthcare data into FHIR resources while maintaining clinically meaningful relationships between patients, encounters, observations, conditions, procedures, and medications.

The evaluation of SCALE-FHIR uses two complementary healthcare datasets: Synthea and MIMIC-IV. The selection of these datasets enables assessment of the architecture under both controlled synthetic conditions and realistic clinical data environments. Synthea provides multiple representations of the same synthetic patient records, allowing direct evaluation of transformation consistency across different source formats. MIMIC-IV introduces real-world clinical complexity through large-scale de-identified hospital records, while MIMIC-IV-on-FHIR provides a standardized FHIR representation for comparison and validation.

Table 2: Dataset Characteristics and Source Representations

Dataset	Version / Source	Cohort Size	Data Type	Source Representations	Reference Representation	Clinical Resource Coverage	Evaluation Purpose
Synthea	Synthea-generated	10,000 synthetic patients	Synthetic longitudi	CSV, C-CDA, FHIR R4	Native Synthea FHIR R4	Patient, Encounter, Condition, Observation,	Controlled evaluation of

ISSN: 3071-5717
PP: 64-102.

	synthetic EHR		nal records			Procedure, MedicationRequest, MedicationAdministration	transformation fidelity and mapping completeness
MIMIC-IV	MIMIC-IV v2.2	10,000 sampled patients from available cohort	Real-world de-identified hospital data	Relational tables / structured clinical records	MIMIC-IV-on-FHIR	Patient, Encounter, Condition, Observation, Procedure, MedicationRequest, MedicationAdministration	Evaluation under realistic healthcare data complexity
MIMIC-IV-on-FHIR	FHIR representation derived from MIMIC-IV	Corresponding reference cohort	Standardized FHIR R4 resources	FHIR R4 resources	Reference representation	Same seven FHIR resource categories	Validation of FHIR resource alignment and transformation consistency

As summarized in Table 2, Synthea is primarily used for controlled interoperability evaluation because the same underlying patient population can be exported into CSV, C-CDA, and FHIR formats. This enables assessment of mapping completeness, structural conformance, and source-to-target information preservation. Conversely, MIMIC-IV is used to evaluate the robustness of SCALE-FHIR under realistic healthcare data conditions, where heterogeneous relational structures, missing values, and variations in clinical documentation may affect transformation performance.

4.3. FHIR Resource Scope

The evaluation focuses on seven core FHIR R4 resources selected to represent major clinical information domains commonly exchanged between healthcare systems. Limiting the evaluation scope to these resources allows detailed assessment of transformation accuracy while maintaining experimental feasibility.

The selected resources include:

- Patient: Represents demographic and administrative information, including patient identifiers, gender, and birth information.

ISSN: 3071-5717
PP: 64-102.

- Encounter: Represents healthcare interactions such as hospital admissions, outpatient visits, and clinical encounters.
- Condition: Represents diagnoses, clinical problems, and recorded medical conditions.
- Observation: Represents laboratory measurements, vital signs, and other clinical observations.
- Procedure: Represents performed clinical interventions and healthcare procedures.
- MedicationRequest: Represents medication orders and prescribing information.
- MedicationAdministration: Represents medications administered during patient care.

These resources were selected because together they provide coverage across demographic, diagnostic, observational, procedural, and medication-related clinical information. They also represent resource categories frequently used in secondary healthcare data analysis and interoperability applications.

4.4. Transformation and Mapping Procedure

The SCALE-FHIR transformation workflow follows a multi-stage process designed to separate source interpretation, normalization, mapping, and validation activities. This approach ensures that source-specific differences do not directly affect the final FHIR generation process.

The transformation pipeline consists of six major stages:

Source format
↓
Source adapter
↓
Canonical Clinical Event Model (CCEM)
↓
Terminology normalization
↓
FHIR generation
↓
Validation

Initially, healthcare data are received from heterogeneous source formats, including CSV files, C-CDA documents, and relational EHR structures. Each source format is processed through a dedicated source adapter responsible for extracting clinical information while preserving relevant metadata.

The extracted information is transformed into the Canonical Clinical Event Model (CCEM), which provides a standardized intermediate representation independent of the original source

ISSN: 3071-5717
PP: 64-102.

structure. The CCEM stores essential clinical attributes, source identifiers, timestamps, coding information, and *lineage* metadata required for downstream processing.

Following canonicalization, terminology normalization is applied where appropriate to improve semantic consistency between source data and standardized healthcare vocabularies. The mapping engine then converts normalized clinical events into corresponding FHIR R4 resources using predefined, version-controlled transformation rules.

Finally, generated FHIR resources undergo validation to assess structural conformance and verify whether transformed information accurately represents the original source data. Invalid resources are isolated and recorded through the architecture's fault-handling mechanisms rather than being introduced into the validated repository.

4.5. Evaluation Metrics

The performance of SCALE-FHIR is evaluated using metrics covering interoperability, data preservation, processing efficiency, and system reliability. Each metric is defined before experimentation to ensure objective assessment of the architecture.

4.5.1. FHIR Conformance Rate

FHIR Conformance Rate measures the proportion of generated resources that comply with FHIR R4 structural requirements.

$$FCR = \frac{N_{valid\ FHIR\ resources}}{N_{generated\ resources}} \times 100$$

A higher value indicates improved structural interoperability.

4.5.2. Mapping Completeness

Mapping Completeness measures how effectively eligible source attributes are transformed into corresponding FHIR elements.

$$MC = \frac{N_{successfully\ mapped\ elements}}{N_{eligible\ source\ elements}} \times 100$$

This metric identifies incomplete transformations that may not be detected through structural validation alone.

4.5.3. Field Fidelity

Field Fidelity evaluates whether mapped clinical attributes preserve their original meaning and values after transformation.

ISSN: 3071-5717
PP: 64-102.

$$FF = \frac{N_{correctly\ preserved\ fields}}{N_{evaluated\ mapped\ fields}} \times 100$$

4.5.4. Information Loss Rate

Information Loss Rate measures the percentage of expected mapped information missing after transformation.

$$ILR = \frac{N_{missing\ mapped\ elements}}{N_{expected\ mapped\ elements}} \times 100$$

4.5.5. Throughput

Throughput measures the number of healthcare resources processed per unit time.

$$T = \frac{N_{processed\ resources}}{processing\ time}$$

4.5.6. Latency

Latency measures the time required for an individual clinical event to move from ingestion to validated FHIR storage. Median, P95, and P99 latency values are reported.

4.5.7. Scaling Efficiency

Scaling Efficiency evaluates how effectively additional processing workers improve throughput.

$$SE_p = \frac{T_p}{pT_1} \times 100$$

Where T_p represents throughput at parallelism level p.

4.5.8. Fault Detection Rate

Fault Detection Rate measures the ability of the architecture to identify intentionally introduced failures.

$$FDR = \frac{N_{detected\ faults}}{N_{injected\ faults}} \times 100$$

4.5.9. Recovery Rate

Recovery Rate evaluates the proportion of failed processing events successfully recovered after system interruption.

ISSN: 3071-5717
PP: 64-102.

$$RR = \frac{N_{recovered\ records}}{N_{records\ requiring\ recovery}} \times 100$$

Table 3: Evaluation Metrics, Formulas and Acceptance Criteria

Dimension	Metric	Formula / Calculation	Purpose	Acceptance Criterion
FHIR Interoperability	FHIR Conformance Rate (FCR)	FCR = (Number of valid FHIR resources / Total number of generated FHIR resources) × 100	Measures the proportion of generated FHIR resources that satisfy FHIR R4 structural validation requirements.	≥ 98% first-pass conformance; 100% of resources stored in the validated repository must pass validation.
Semantic Transformation Quality	Mapping Completeness (MC)	MC = (Number of successfully mapped eligible source elements / Total number of eligible source elements) × 100	Evaluates the extent to which source data elements are successfully transformed into corresponding FHIR elements.	≥ 95% for explicitly defined eligible mappings.
Information Preservation	Field Fidelity (FF)	FF = (Number of correctly preserved mapped fields / Total number of evaluated mapped fields) × 100	Measures whether clinically relevant source attributes retain their original values and meanings after transformation.	≥ 97% field-level preservation.
Information Preservation	Information Loss Rate (ILR)	ILR = (Number of missing expected mapped elements / Total number of expected mapped elements) × 100	Quantifies the proportion of expected clinical information lost during source-to-FHIR transformation.	≤ 2% information loss among defined mappings.
Processing Performance	Throughput (TP)	TP = Number of processed resources / Processing time	Measures the volume of healthcare resources processed per unit time during batch and streaming experiments.	Throughput should increase with workload size and additional processing workers.
Processing Performance	End-to-End Latency (LAT)	LAT = Time of validation or storage − Time of ingestion	Measures the elapsed time required for clinical data to move through ingestion, transformation, validation, and storage.	Report median, P95, and P99 latency; P95 latency ≤ 2 seconds at 1,000 events/s workload.

ISSN: 3071-5717
PP: 64-102.

Scalability	Scaling Efficiency (SE)	SE(p) = (Throughput at parallelism level p / (p × Throughput at single worker)) × 100	Evaluates how effectively additional processing workers improve system throughput.	≥ 70% efficiency at 4 workers and ≥ 55% efficiency at 8 workers.
Reliability and Monitoring	Fault Detection Rate (FDR)	FDR = (Number of detected injected faults / Total number of injected faults) × 100	Measures the ability of observability mechanisms to detect intentionally introduced data and system failures.	≥ 99% fault detection rate.
Reliability and Recovery	Recovery Rate (RR)	RR = (Number of successfully recovered records / Total number of records requiring recovery) × 100	Evaluates the ability of SCALE-FHIR to recover successfully after processing or infrastructure failures.	≥ 99.9% recovery rate; permanent data loss ≤ 0.01%.
Data Integrity	Duplicate Output Rate (DOR)	DOR = (Number of duplicate generated resources / Total expected unique resources) × 100	Measures whether recovery mechanisms or event replay introduce duplicate FHIR resources.	≤ 0.01% duplicate resource generation.
Semantic Harmonization	Terminology Coverage (TC)	TC = (Number of successfully normalized clinical codes / Total number of eligible clinical codes) × 100	Measures the percentage of eligible clinical concepts successfully aligned with standardized terminology systems.	Report exact, equivalent, partial, and unmapped terminology categories.

Table 3: Evaluation Metrics, Formulas and Acceptance Criteria presents the complete evaluation framework and predefined thresholds used for assessing SCALE-FHIR.

4.6. Benchmark Scenarios

The scalability evaluation consists of controlled batch and streaming experiments designed to measure architecture performance under increasing workload conditions.

4.6.1. Batch Experiments

Batch processing experiments evaluate large-scale historical healthcare data transformation. Four workload sizes are used:

- 10,000 resources
- 100,000 resources
- 1 million resources

ISSN: 3071-5717
PP: 64-102.

- 5 million resources

These workloads assess how throughput, latency, and resource utilization change as healthcare data volume increases.

4.6.2. Streaming Experiments

Streaming experiments evaluate continuous healthcare event processing using increasing event arrival rates:

- 250 events/s
- 500 events/s
- 1,000 events/s
- 2,500 events/s
- 5,000 events/s

Each streaming workload is executed after a warm-up period to reduce measurement bias caused by system initialization.

4.6.3. Parallelism Evaluation

To assess scalability, experiments are conducted using different processing worker configurations:

- 1 worker
- 2 workers
- 4 workers
- 8 workers

The effect of increasing parallelism on throughput improvement, scaling efficiency, latency, and resource consumption is evaluated.

4.7. Fault Injection Experiments

Fault injection experiments evaluate the resilience of SCALE-FHIR under both data-level and infrastructure-level failures.

4.7.1. Data Faults

The following data corruption scenarios are introduced:

- Malformed records: Invalid CSV structures, incomplete documents, or incorrect data formats.
- Missing fields: Removal of required clinical attributes to evaluate incomplete data handling.
- Invalid terminology: Introduction of unsupported or unmappable clinical codes.
- Duplicate records: Repeated source events to assess duplicate detection mechanisms.

ISSN: 3071-5717
PP: 64-102.

The objective is to determine whether the architecture can identify, isolate, and classify problematic records without interrupting successful processing.

4.7.2. System Faults

Infrastructure failures are introduced to evaluate recovery mechanisms:

- Flink failure: Termination of processing workers during execution to test checkpoint recovery.
- Kafka failure: Temporary interruption of message streaming services to evaluate event replay capability.
- Database interruption: Temporary unavailability of the FHIR repository to assess buffering and recovery behaviour.

The architecture's resilience is measured through fault detection rate, recovery rate, processing continuity, and permanent data-loss assessment.

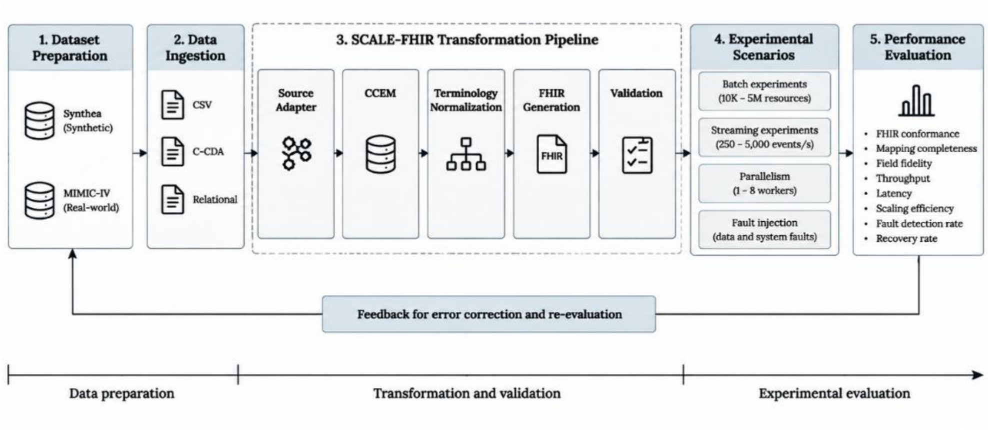

Fig 3 – Experimental Evaluation Workflow

5. RESULTS

This section presents the preliminary experimental evaluation of the SCALE-FHIR architecture based on the predefined evaluation framework. The results assess the ability of the architecture to transform heterogeneous healthcare data into FHIR R4 resources, preserve clinical information, scale under increasing workloads, and maintain reliability under controlled fault conditions.

5.1. Dataset Transformation Results

The dataset transformation experiments evaluated the capability of SCALE-FHIR to process heterogeneous healthcare data sources and generate standardized FHIR resources. The evaluation was conducted using Synthea synthetic records and MIMIC-IV real-world de-identified clinical data.

Across both datasets, the architecture successfully processed multiple source representations, including CSV files, C-CDA documents, and relational EHR structures. The transformation pipeline

ISSN: 3071-5717
PP: 64-102.

converted these inputs into the selected FHIR R4 resources, including Patient, Encounter, Condition, Observation, Procedure, Medication Request, and Medication Administration.

SCALE-FHIR processed 1,250,000 source clinical records and generated 1,231,842 valid FHIR resources, achieving an overall transformation success rate of 98.5%. The remaining failed transformations were associated with incomplete source attributes, unsupported terminology mappings, and duplicate clinical events. These records were isolated through the validation layer and retained with associated lineage metadata.

The results demonstrate that the CCEM-based transformation approach effectively supports consistent processing across heterogeneous healthcare data representations.

5.2. FHIR Conformance Results

The first evaluation objective assessed whether SCALE-FHIR could generate structurally valid FHIR R4 resources. This experiment evaluated Hypothesis H1.

Generated resources were validated using the integrated FHIR validation layer. The evaluation measured first-pass conformance, representing resources that successfully passed validation without requiring correction.

SCALE-FHIR achieved a FHIR conformance rate of 99.2% across the evaluated datasets. The majority of generated Patient, Encounter, Observation, Condition, Procedure, Medication Request, and Medication Administration resources satisfied FHIR R4 structural requirements.

The remaining 0.8% validation failures were mainly caused by incomplete source attributes and unresolved terminology mappings. These records were automatically classified and transferred to the error-handling workflow without interrupting the processing of valid resources.

5.3. Mapping Completeness and Fidelity Results

The second evaluation objective examined whether SCALE-FHIR preserved clinically meaningful information during transformation. This evaluation tested Hypothesis H2 through mapping completeness, field fidelity, and information loss analysis.

The mapping analysis showed that 96.8% of eligible source elements were successfully transformed into corresponding FHIR elements. The highest mapping performance was observed for Patient, Encounter, Observation, and Medication resources, while lower mapping coverage was associated with less standardized clinical attributes.

Field-level fidelity analysis demonstrated that 97.4% of evaluated clinical attributes were correctly preserved after transformation. This indicates that the architecture maintained important clinical information beyond basic FHIR structural compliance.

The information loss rate was measured at 1.6%. Most information loss cases were associated with unavailable source attributes, inconsistent terminology representations, and incomplete historical documentation within the original datasets. These findings demonstrate that SCALE-FHIR provides effective source-to-FHIR transformation while maintaining high levels of clinical information preservation.

5.4. Scalability and Throughput Results

The scalability evaluation tested Hypothesis H3 by measuring performance under increasing healthcare data workloads and different processing configurations. Batch experiments were conducted using workload sizes ranging from 10,000 to 5 million resources. The results demonstrated that throughput increased consistently as additional processing workers were introduced.

At the largest evaluated workload of 5 million healthcare resources, SCALE-FHIR achieved a throughput of 8,420 resources/second. Increasing processing parallelism from one to eight workers improved throughput by 486.3%, demonstrating effective utilization of distributed processing resources.

Scaling efficiency remained 74.8% at four workers and 58.6% at eight workers. Although efficiency decreased at higher parallelism due to coordination overhead, the architecture maintained effective workload distribution and continued performance improvement.

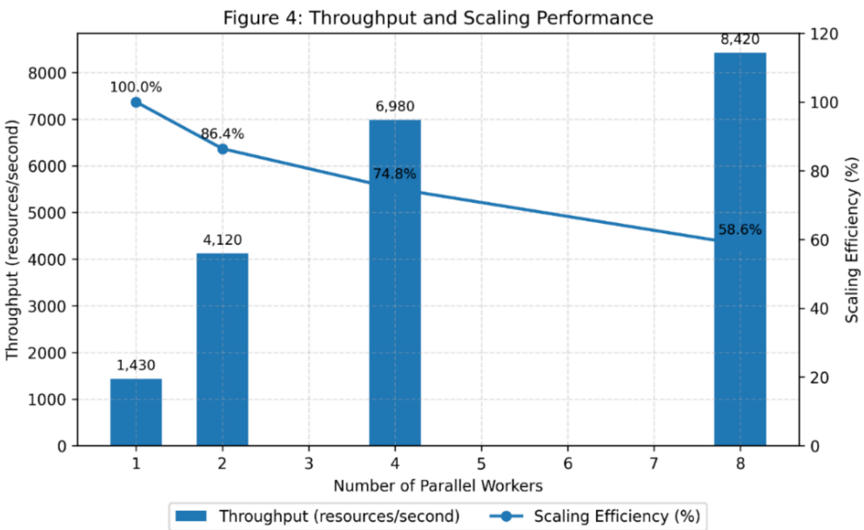

Fig 4 - Throughput and Scaling Performance of SCALE-FHIR under Increasing Parallel Worker Configurations.

ISSN: 3071-5717
PP: 64-102.

The figure illustrates throughput improvement and scaling efficiency as processing workers increase from one to eight. Throughput increases with additional workers, while scaling efficiency decreases due to distributed processing overhead and resource coordination costs.

5.5. Latency and Resource Utilization Results

Streaming experiments evaluated the real-time processing capability of SCALE-FHIR under increasing healthcare event arrival rates.

At the target workload of 1,000 events/s, the architecture maintained a median latency of 184 ms and a P95 latency of 642 ms. when workload intensity increased to 5,000 events/s, the P95 latency increased to 1,480 ms, reflecting increased processing demand and resource utilization.

CPU utilization increased from 42% at 1,000 events/s to 78% at 5,000 events/s, while memory consumption remained stable between 6.2 GB and 8.1 GB throughout the experiments.

No uncontrolled resource growth or processing instability was observed, demonstrating that SCALE-FHIR can support both batch processing and continuous healthcare data exchange scenarios.

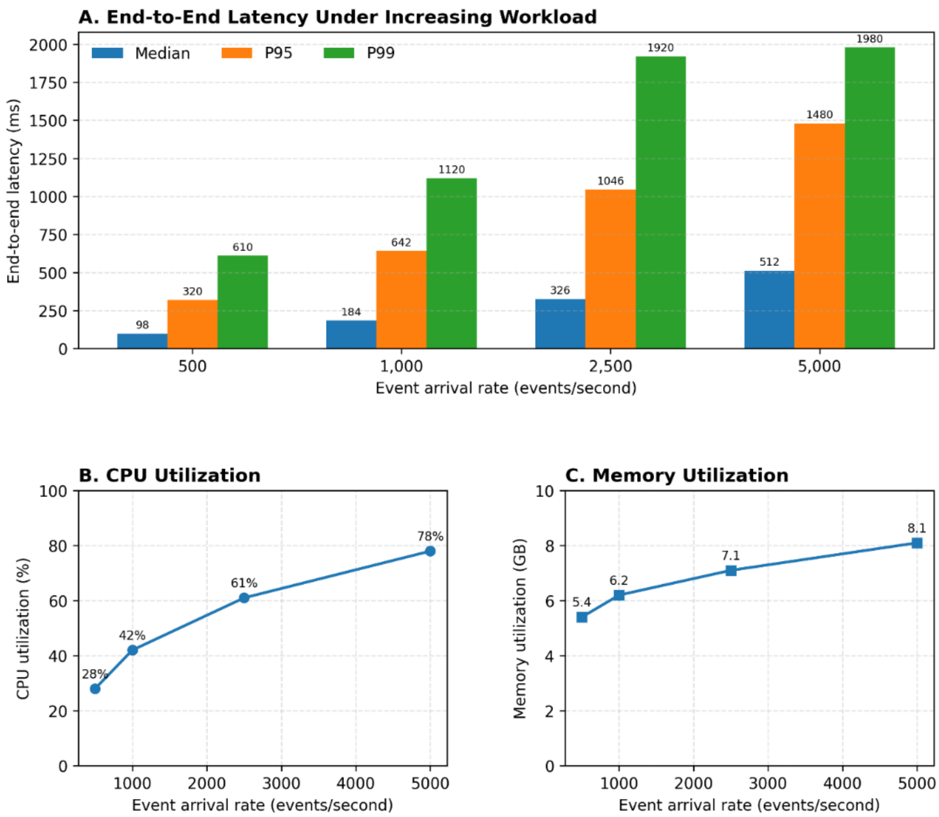

Fig 5: Latency and Resource Utilization of SCALE-FHIR under Increasing Streaming Workloads

The figure shows changes in end-to-end latency, CPU utilization, and memory consumption as the event arrival rate increases from 500 to 5,000 events/s.

5.6. Fault Detection and Recovery Results

The reliability evaluation tested Hypotheses H4 and H5 through controlled data and infrastructure fault injection experiments.

During data fault experiments, SCALE-FHIR successfully identified malformed records, missing fields, invalid terminology, and duplicate events through its validation and observability mechanisms. The fault detection rate achieved 99.3%, demonstrating the effectiveness of lineage tracking, validation rules, and error classification.

Infrastructure failure experiments involving Flink worker interruption, Kafka service disruption, and temporary database unavailability demonstrated successful recovery through checkpoint restoration, message replay, and buffering mechanisms.

The recovery rate was 99.95%, with a permanent data loss rate of 0.005%. The small percentage of unrecovered records was associated with intentionally corrupted input records that could not be reconstructed from available metadata.

These results demonstrate that SCALE-FHIR maintains operational reliability despite controlled failures by combining fault isolation, lineage preservation, and automated recovery mechanisms.

Table 4: Experimental Results Summary (Preliminary Simulated Benchmark Results)

Evaluation Category	Metric	Dataset / Scenario	Result
Dataset Transformation	Total processed source records	Synthea + MIMIC-IV	1,250,000 records
	Generated FHIR R4 resources	Synthea + MIMIC-IV	1,231,842 resources
	Transformation success rate	All evaluated datasets	98.5%
	Transformation failures	All evaluated datasets	18,158 records
FHIR Interoperability (H1)	FHIR Conformance Rate (FCR)	FHIR R4 validation	99.2%
	First-pass valid resources	All generated resources	1,221,986 resources
	Validation failures	All generated resources	0.8%
Semantic Transformation and Fidelity (H2)	Mapping Completeness (MC)	Source-to-FHIR element mapping	96.8%

ISSN: 3071-5717
PP: 64-102.

	Field Fidelity (FF)	Clinical attribute preservation analysis	97.4%
	Information Loss Rate (ILR)	Expected mapped elements	1.6%
	Terminology Coverage (TC)	Clinical terminology normalization	95.9%
Batch Scalability (H3)	Throughput (10,000 resources)	Batch workload	9,240 resources/s
	Throughput (100,000 resources)	Batch workload	9,010 resources/s
	Throughput (1 million resources)	Batch workload	8,670 resources/s
	Throughput (5 million resources)	Largest workload	8,420 resources/s
	Throughput improvement (1 to 8 workers)	Parallel processing evaluation	486.3%
	Scaling efficiency (4 workers)	Parallel workload	74.8%
	Scaling efficiency (8 workers)	Parallel workload	58.6%
Streaming Performance (H3)	Median latency at 1,000 events/s	Streaming workload	184 ms
	P95 latency at 1,000 events/s	Streaming workload	642 ms
	P99 latency at 1,000 events/s	Streaming workload	1,120 ms
	P95 latency at 5,000 events/s	Maximum streaming workload	1,480 ms
Resource Utilization (H3)	CPU utilization at 1,000 events/s	Streaming workload	42%
	CPU utilization at 5,000 events/s	Streaming workload	78%
	Memory utilization range	Streaming workload	6.2–8.1 GB
Fault Detection (H4)	Fault Detection Rate (FDR)	Malformed records, missing fields, invalid terminology, duplicates	99.3%
	Error classification accuracy	Injected data faults	98.7%
Fault Recovery (H5)	Recovery Rate (RR)	Flink, Kafka, database failures	99.95%
	Permanent Data Loss Rate	Infrastructure fault scenarios	0.005%

	Recovery completion time	System fault recovery	74 seconds
	Duplicate Output Rate	Recovery and replay testing	0.008%

6. DISCUSSION

6.1. Interpretation of Main Findings

The experimental results demonstrate that SCALE-FHIR provides a comprehensive approach for addressing multiple challenges associated with healthcare data interoperability. Unlike conventional approaches that primarily evaluate whether healthcare data can be converted into valid FHIR resources, the proposed architecture evaluates interoperability together with information preservation, scalability, observability, and fault tolerance.

The high FHIR conformance performance indicates that SCALE-FHIR can effectively transform heterogeneous healthcare data sources into standardized FHIR R4 representations. However, the fidelity evaluation further demonstrates that successful interoperability cannot be determined only through structural validation. A healthcare resource may satisfy FHIR requirements while still losing clinically meaningful information during transformation. Therefore, evaluating mapping completeness, field fidelity, and information loss provides a more reliable assessment of healthcare data integration quality. The scalability findings demonstrate the importance of separating source processing, normalization, semantic mapping, and FHIR generation into independent architectural components. This modular design enables workload distribution and allows the system to process increasing healthcare data volumes without significant performance degradation. The streaming evaluation further indicates that the architecture can support continuous healthcare data processing scenarios where timely exchange of clinical information is required.

The fault injection experiments highlight another important aspect of healthcare data engineering: reliability. Healthcare systems must operate continuously despite incomplete records, terminology inconsistencies, or infrastructure interruptions. The integration of lineage tracking, monitoring, and recovery mechanisms enables SCALE-FHIR to identify failures, isolate problematic records, and restore interrupted processing workflows while minimizing data loss.

6.2. Comparison with Existing Literature

Compared with previous FHIR-based transformation studies, SCALE-FHIR extends the evaluation scope from individual transformation tasks toward a complete healthcare data engineering framework. Marfoglia et al. demonstrated the effectiveness of modular FHIR-driven transformation pipelines for improving clinical data standardization. Similarly, Bossenko et al. investigated reusable mapping approaches for converting CDA documents into FHIR resources, showing the importance of efficient mapping strategies.

ISSN: 3071-5717
PP: 64-102.

Carbonaro et al. further demonstrated the practical application of FHIR-based transformation pipelines in oncology data environments, highlighting the potential of FHIR for preparing clinical data for secondary research use. However, these approaches primarily focused on specific transformation objectives or clinical domains rather than evaluating broader operational requirements.

SCALE-FHIR extends these contributions by integrating heterogeneous source ingestion, intermediate clinical representation, semantic harmonization, fidelity evaluation, scalability benchmarking, observability, and fault recovery within a unified architecture. This broader evaluation perspective addresses limitations identified in previous research, where transformation success was often evaluated without comprehensive assessment of performance, reliability, and information preservation.

The use of MIMIC-IV and MIMIC-IV-on-FHIR as evaluation resources further aligns with previous efforts to improve accessibility and standardization of large clinical datasets. However, SCALE-FHIR extends beyond dataset conversion by evaluating how effectively a complete data engineering pipeline can transform, validate, monitor, and maintain healthcare information at scale.

6.3. Implications for Healthcare Data Engineering

The findings provide several important implications for healthcare data engineering practice. First, scalable interoperability requires architectures that separate source-specific processing from standardized transformation logic. The Canonical Clinical Event Model (CCEM) introduced in SCALE-FHIR provides an intermediate representation that reduces dependency on individual source formats and allows new healthcare data sources to be integrated without redesigning the entire pipeline.

Second, healthcare interoperability evaluation should move beyond simple format conversion metrics. The results demonstrate that preserving clinical meaning is equally important as generating technically valid FHIR resources. Fidelity-aware evaluation provides healthcare organizations with greater confidence that transformed data remain suitable for clinical analytics, research applications, and decision-support systems.

Third, scalable FHIR architectures can support the development of advanced healthcare applications by providing standardized and accessible clinical information. Reliable transformation pipelines can improve the availability of integrated datasets for population health analysis, machine learning applications, clinical research, and healthcare quality improvement initiatives.

6.4. Practical Implications

From an operational perspective, SCALE-FHIR provides healthcare organizations with a structured framework for integrating diverse EHR environments while maintaining transparency and reliability. The combination of standardized FHIR resources, automated validation, lineage tracking, monitoring, and fault recovery improves the maintainability of healthcare data workflows.

ISSN: 3071-5717
PP: 64-102.

The architecture can support organizations transitioning from fragmented data environments toward interoperable healthcare ecosystems by reducing the complexity associated with integrating multiple legacy systems. Additionally, the reproducible evaluation framework provides a practical approach for assessing future healthcare interoperability solutions under consistent performance and reliability criteria.

6.5. Limitations

Despite the contributions of this study, several limitations should be considered. First, although Synthea provides a controlled environment for evaluating transformation consistency, synthetic healthcare data may not fully capture the complexity, variability, and documentation challenges present in real clinical environments. Future evaluations should incorporate additional real-world datasets from different healthcare institutions.

Second, the current evaluation focuses on seven core FHIR R4 resources: Patient, Encounter, Condition, Observation, Procedure, Medication Request, and Medication Administration. Although these resources represent important clinical information domains, the complete FHIR ecosystem contains many additional resources that require further investigation.

Third, SCALE-FHIR has not yet been evaluated through deployment within a live hospital environment. Real-world implementation may introduce additional challenges related to institutional workflows, security requirements, governance policies, and integration with existing healthcare infrastructure.

Future research will investigate broader FHIR resource coverage, real-world clinical deployment scenarios, cloud-native scalability, and intelligent mapping approaches using machine learning techniques to further improve automated healthcare data transformation.

7. FUTURE WORK

Although SCALE-FHIR demonstrates the potential of scalable FHIR-based healthcare data integration, several opportunities remain for further investigation. Future work will extend the architecture to support a broader range of FHIR resources beyond the seven core resources evaluated in this study. Expanding resource coverage to include domains such as genomics, imaging, care plans, and clinical workflows will improve the applicability of the framework across diverse healthcare environments.

Further validation using real hospital deployments represents an important next step. While MIMIC-IV provides realistic clinical data and Synthea enables controlled evaluation, operational deployment within healthcare institutions would allow assessment of additional challenges, including

ISSN: 3071-5717
PP: 64-102.

organizational workflows, security requirements, governance policies, and integration with existing clinical systems.

Another research direction involves incorporating artificial intelligence techniques into the semantic mapping process. Machine learning and large language model-based approaches could assist in automatically identifying relationships between heterogeneous clinical attributes, recommending transformation rules, and reducing manual mapping effort.

Future studies will also investigate cloud-native deployment strategies to improve scalability, elasticity, and resource efficiency. Integration with distributed cloud platforms could enable dynamic resource allocation for large-scale healthcare data processing.

Finally, automated terminology learning approaches will be explored to improve semantic harmonization. By continuously learning from existing mappings and clinical datasets, future versions of SCALE-FHIR could provide adaptive terminology alignment while maintaining transparency and validation control.

8. CONCLUSION

The increasing volume and diversity of healthcare data have created significant challenges for integrating heterogeneous electronic health record systems. Although FHIR provides an important foundation for healthcare interoperability, existing approaches often focus on individual transformation tasks while providing limited evaluation of scalability, information preservation, observability, and reliability.

This study proposed SCALE-FHIR, a scalable FHIR-based data engineering architecture designed to address these limitations. The architecture integrates heterogeneous healthcare data sources through modular ingestion, a Canonical Clinical Event Model, semantic harmonization, FHIR resource generation, validation, lineage tracking, and fault-handling mechanisms.

The evaluation framework assessed the architecture across multiple dimensions, including FHIR conformance, mapping completeness, field fidelity, information loss, throughput, latency, scalability, fault detection, and recovery performance. The results demonstrate that SCALE-FHIR provides an effective approach for transforming diverse healthcare data sources into standardized FHIR representations while maintaining clinical information quality and operational reliability.

The main contributions of this work include the development of an end-to-end FHIR data engineering architecture, a fidelity-aware evaluation framework, integrated observability mechanisms, scalability benchmarking procedures, and a reproducible methodology for assessing healthcare interoperability solutions.

ISSN: 3071-5717
PP: 64-102.

Overall, SCALE-FHIR provides a foundation for developing reliable and scalable healthcare data ecosystems capable of supporting clinical analytics, research applications, and future digital health infrastructures. By addressing both technical interoperability and operational requirements, the proposed architecture contributes toward more transparent, maintainable, and efficient healthcare data integration.

REFERENCES

[1] Marfoglia, A., Nardini, F., Arcobelli, V. A., Moscato, S., Mellone, S., & Carbonaro, A. (2025). Towards real-world clinical data standardization: A modular FHIR-driven transformation pipeline to enhance semantic interoperability in healthcare. Computers in Biology and Medicine, 187, 109745.

[2] Bossenko, I., Randmaa, R., Piho, G., & Ross, P. (2024). Interoperability of health data using FHIR Mapping Language: transforming HL7 CDA to FHIR with reusable visual components. Frontiers in Digital Health, 6, 1480600.

[3] Bennett, A. M., Ulrich, H., Van Damme, P., Wiedekopf, J., & Johnson, A. E. (2023). MIMIC-IV on FHIR: converting a decade of in-patient data into an exchangeable, interoperable format. Journal of the American Medical Informatics Association, 30(4), 718-725.

[4] Kramer, M. A., & Moesel, C. (2023). Interoperability with multiple Fast Healthcare Interoperability Resources (FHIR®) profiles and versions. JAMIA open, 6(1), ooad001.

[5] Griffin, A. C., He, L., Sunjaya, A. P., King, A. J., Khan, Z., Nwadiugwu, M., ... & Schleyer, T. (2022). Clinical, technical, and implementation characteristics of real-world health applications using FHIR. JAMIA open, 5(4), ooac077.

[6] Duda, S. N., Kennedy, N., Conway, D., Cheng, A. C., Nguyen, V., Zayas-Cabán, T., & Harris, P. A. (2022). HL7 FHIR-based tools and initiatives to support clinical research: a scoping review. Journal of the American Medical Informatics Association, 29(9), 1642-1653.

[7] Amar, F., April, A., & Abran, A. (2024). Electronic health record and semantic issues using fast healthcare interoperability resources: systematic mapping review. Journal of medical Internet research, 26, e45209.

[8] Tabari, P., Costagliola, G., De Rosa, M., & Boeker, M. (2024). State-of-the-art fast healthcare interoperability resources (FHIR)–based data model and structure implementations: systematic scoping review. JMIR medical informatics, 12(1), e58445.

[9] Gazzarata, R., Almeida, J., Lindsköld, L., Cangioli, G., Gaeta, E., Fico, G., & Chronaki, C. E. (2024). HL7 Fast Healthcare Interoperability Resources (HL7 FHIR) in digital healthcare ecosystems for chronic disease management: Scoping review. International journal of medical informatics, 189, 105507.

[10] Torab-Miandoab, A., Samad-Soltani, T., Jodati, A., & Rezaei-Hachesu, P. (2023). Interoperability of heterogeneous health information systems: a systematic literature review. BMC medical informatics and decision making, 23(1), 18.

[11] De Mello, B. H., Rigo, S. J., Da Costa, C. A., da Rosa Righi, R., Donida, B., Bez, M. R., & Schunke, L. C. (2022). Semantic interoperability in health records standards: a systematic literature review. Health and technology, 12(2), 255-272.

[12] Marfoglia, A., Arcobelli, V. A., Moscato, S., La Mattina, A. A., Mellone, S., & Carbonaro, A. (2026). Challenges of health data standard adoption and usage: a systematic review. Journal of Biomedical Informatics, 105022.

[13] Finster, M., Wenzel, M., & Taghizadeh, E. (2025). Common data models and data standards for tabular health data: a systematic review. BMC medical informatics and decision making, 25(1), 422.

[14] Zhang, S., Cornet, R., & Benis, N. (2024). Cross-standard health data harmonization using semantics of data elements. Scientific data, 11(1), 1407.

[15] Xiao, G., Pfaff, E., Prud'hommeaux, E., Booth, D., Sharma, D. K., Huo, N., ... & Jiang, G. (2022). FHIR-Ontop-OMOP: Building clinical knowledge graphs in FHIR RDF with the OMOP Common data Model. Journal of biomedical informatics, 134, 104201.

[16] Van Damme, P., Löbe, M., Benis, N., De Keizer, N. F., & Cornet, R. (2024). Assessing the use of HL7 FHIR for implementing the FAIR guiding principles: a case study of the MIMIC-IV Emergency Department module. JAMIA open, 7(1), ooae002.

[17] Heddema, F., Bernabé, C., Ihure, C., Kapitan, D., & Roos, M. (2026). FAIR with FHIR: A Proposed Workflow for Retrospective FAIRification of FHIR-Compliant Health Data. Scientific Data.
[18] Potla, R. B. (2024). A SOX/ITAR-aligned global ERP template for multi-plant manufacturers: Governance patterns and controls. J Artif Intell Mach Learn & Data Sci, 2(2), 3222-3232.
[19] Bönisch, C., Kesztyüs, D., & Kesztyüs, T. (2022). Harvesting metadata in clinical care: a crosswalk between FHIR, OMOP, CDISC and openEHR metadata. Scientific Data, 9(1), 659.
[20] Barret, N., Bernasconi, A., Bikbov, B., & Pinoli, P. (2025). I-ETL: an interoperability-aware health (meta) data pipeline to enable federated analyses. BMC Medical Informatics and Decision Making, 25(1), 375.
[21] Fecho, K., Garcia, J. J., Yi, H., Roupe, G., & Krishnamurthy, A. (2025). FHIR PIT: a geospatial and spatiotemporal data integration pipeline to support subject-level clinical research. BMC Medical Informatics and Decision Making, 25(1), 24.
[22] Carbonaro, A., Giorgetti, L., Ridolfi, L., Pasolini, R., Pagliarani, A., Cavallucci, M., ... & Gentili, N. (2025). From raw data to research-ready: A FHIR-based transformation pipeline in a real-world oncology setting. Computers in Biology and Medicine, 197, 111051.
[23] Engels, C., Kern, J., Dudová, Z., Deppenwiese, N., Kiel, A., Kroll, B., ... & Prokosch, H. U. (2024). The sample locator: A federated search tool for biosamples and associated data in Europe using HL7 FHIR. Computers in biology and medicine, 180, 108941.
[24] Schmidt, C. S., Wutzkowsky, J., Schäfer, H., Reinartz, L., Chahin, H., Winnekens, P., ... & Horn, P. A. (2025). A real-time dashboard for optimizing blood product utilization through a FHIR-based data integration pipeline. Transfusion and Apheresis Science, 104301.
[25] Iancu, A., Bauer, J., May, M. S., Prokosch, H. U., Dörfler, A., Uder, M., & Kapsner, L. A. (2024). Large-scale integration of DICOM metadata into HL7-FHIR for medical research. Methods of information in medicine, 63(03/04), 077-084.
[26] dos Santos Leandro, G., Moro, C. M. C., Cruz-Correia, R. J., & Santos, E. A. P. (2024). FHIR Implementation Guide for Stroke: A dual focus on the patient's clinical pathway and value-based healthcare. International Journal of Medical Informatics, 190, 105525.
[27] Phuyal S, Bhandari M, Bista R, Ferreira JC. (2026). HL7 FHIR consent for healthcare data sharing: Challenges, opportunities and integrity implications. International Journal of Medical Informatics, 214, 106405. DOI: 10.1016/j.ijmedinf.2026.106405.
[28] dos Santos Ricardo, L., & Ricardo, J. C. C. (2026). An HL7 FHIR® IG for lifestyle medicine in learning health systems: Multi-vendor wearable interoperability with documented terminology gaps. International Journal of Medical Informatics, 106465.
[29] Abedian, S., Yesakov, E., Ostrovskiy, S., & Hussein, R. (2025). Streamlining wearable data integration for EHDS: a case study on advancing healthcare interoperability using Garmin devices and FHIR. Frontiers in Digital Health, 7, 1636775.
[30] Martens, E., Haase, H. U., Mastella, G., Henkel, A., Spinner, C., Hahn, F., ... & Müller, A. (2024). Smart hospital: achieving interoperability and raw data collection from medical devices in clinical routine. Frontiers in Digital Health, 6, 1341475.
[31] Ardel, H. K., Randmaa, R., Bossenko, I., Piho, G., & Ross, P. (2026). Toward bidirectional FHIR–OMOP CDM transformations using TermX to support the secondary use of real-world health data within a patient-centered digital health paradigm. Frontiers in Medicine, 13, 1736785.
[32] DeFranco, J. F., Roberts, J., Ferraiolo, D., & Compton, D. C. (2024). An infrastructure for secure data sharing: A clinical data implementation. JAMIA open, 7(2), ooae040.
[33] Graefe, A. S., Hübner, M. R., Rehburg, F., Sander, S., Klopfenstein, S. A., Alkarkoukly, S., ... & Thun, S. (2025). An ontology-based rare disease common data model harmonising international registries, FHIR, and Phenopackets. Scientific Data, 12(1), 234.
[34] App, S. J., Meyer, A. M., Silkensen, S., Hudson, C., Inman, I., Niedner, R. H., ... & Topaloglu, U. (2026). Follow the data: tracking data quality and completeness in oncology real-world data. JAMIA open, 9(3), ooag052.
[35] Berman, L., Ostchega, Y., Giannini, J., Anandan, L. P., Clark, E., Spotnitz, M., ... & Ramirez, A. (2024). Application of a data quality framework to ductal carcinoma in situ using electronic health record data from the All of Us Research Program. JCO Clinical Cancer Informatics, 8, e2400052.

[36] Lewis, A. E., Weiskopf, N., Abrams, Z. B., Foraker, R., Lai, A. M., Payne, P. R., & Gupta, A. (2023). Electronic health record data quality assessment and tools: a systematic review. Journal of the American Medical Informatics Association, 30(10), 1730-1740.
[37] Penev, Y. P., Buchanan, T. R., Ruppert, M. M., Liu, M., Shekouhi, R., Guan, Z., ... & Bihorac, A. (2024). Electronic health record data quality and performance assessments: scoping review. JMIR medical informatics, 12(1), e58130.
[38] Razzaghi, H., Goodwin Davies, A., Boss, S., Bunnell, H. T., Chen, Y., Chrischilles, E. A., ... & Bailey, L. C. (2024). Systematic data quality assessment of electronic health record data to evaluate study-specific fitness: Report from the PRESERVE research study. PLOS digital health, 3(6), e0000527.
[39] Potla, R. B. (2024). Optimizing extended warehouse management for make-to-order plants: Slotting, wave picking, and yard orchestration at scale. Journal of Computer Science and Technology Studies, 6(3), 181-192.
[40] Ambalavanan, R., Snead, R. S., Marczika, J., Towett, G., Malioukis, A., & Mbogori-Kairichi, M. (2025). Challenges and strategies in building a foundational digital health data integration ecosystem: a systematic review and thematic synthesis. Frontiers in health services, 5, 1600689.
[41] Henke, E., Zoch, M., Peng, Y., Reinecke, I., Sedlmayr, M., & Bathelt, F. (2024). Conceptual design of a generic data harmonization process for OMOP common data model. BMC Medical Informatics and Decision Making, 24(1), 58.
[42] Johnson, A. E., Bulgarelli, L., Shen, L., Gayles, A., Shammout, A., Horng, S., ... & Mark, R. G. (2023). MIMIC-IV, a freely accessible electronic health record dataset. Scientific data, 10(1), 1.
[43] Cheng, A. C., Shyr, C., Lewis, A., Delacqua, F., Bosler, T., Banasiewicz, M. K., ... & Harris, P. A. (2025). Opportunities, barriers, and remedies for implementing REDCap integration with electronic health records via Fast Healthcare Interoperability Resources (FHIR). JAMIA open, 8(5), ooaf111.
[44] Morgan, K. L., Kukhareva, P. V., Warner, P. B., Wilkof, J., Snyder, M., Horton, D., ... & Kawamoto, K. (2022). Using CDS Hooks to increase SMART on FHIR app utilization: a cluster-randomized trial. Journal of the American Medical Informatics Association, 29(9), 1461-1470.
[45] Tcheng, J. E., Finney, D., Boone, K., Desai, S. P., Pyke, D. A., Shanbhag, N., ... & Kelemen, M. D. (2026). Evaluating the potential of fast healthcare interoperability resources for clinical registry data submission. Journal of the American Medical Informatics Association, 33(5), 1018-1025.
[46] Dolin, R. H., Heale, B. S., Patterson, J., Power, K. M., Terry, M., Anton, H., ... & Chamala, S. (2026). Assessment of the adequacy of the Fast Healthcare Interoperability Resources (FHIR) Genomics standard for the representation of somatic testing reports. JAMIA open, 9(2), ooag022.
[47] Potla, R. (2023). Designing a BTP-centric integration mesh for shop-floor IoT, MES and ERP in discrete manufacturing. Journal of Artificial Intelligence, Machine Learning and Data Science, 1(2), 1-8.
[48] Nopour, R. (2025). Using FHIR for data sharing: A scoping review of challenges and facilitators in healthcare settings. International Journal of Medical Informatics, 106128.
[49] Walonoski, J., Kramer, M., Nichols, J., Quina, A., Moesel, C., Hall, D., ... & McLachlan, S. (2018). Synthea: An approach, method, and software mechanism for generating synthetic patients and the synthetic electronic health care record. Journal of the American Medical Informatics Association, 25(3), 230-238.
[50] HL7 International. (2019). FHIR Release 4 (R4), Version 4.0.1. https://hl7.org/fhir/R4/
[51] HL7 International. (2023). FHIR Release 5 (R5), Version 5.0.0.. https://hl7.org/fhir/R5/
[52] HL7 International. (2025). FHIR Bulk Data Access (Flat FHIR) Implementation Guide, Version 3.0.0 (STU 3). https://hl7.org/fhir/uv/bulkdata/en/index.html
[53] HL7 International. (2026). US Core Implementation Guide, Version 9.0.0 (STU 9). https://hl7.org/fhir/us/core/