



Unified Data Foundations and Graph-Based Temporal Modeling for Healthcare Service Operations: Architectural Patterns Supporting Omnichannel Pharmacy and Payer Workflows

Mohammad Jakeer Mehathar*

Independent Researcher, USA

* Corresponding Author Email: jakeer.mehathar@gmail.com - ORCID: 0000-0002-0247-0030

Article Info:

DOI: 10.22399/ijcesen.4960

Received : 03 January 2026

Revised : 20 February 2026

Accepted : 22 February 2026

Keywords

Pharmacy Service Operations,
Unified Data Foundations,
Event-Sourced Platforms,
Graph-Based Temporal Modeling,
Saga Coordination Patterns

Abstract:

Pharmacy service centers and payer operations demand instantaneous access to prescription status, benefit determinations, authorization outcomes, and fulfillment tracking data. Organizational realities frequently involve disconnected platforms handling claims adjudication, eligibility verification, authorization management, dispensing logistics, and relationship tracking separately. Staff members reconcile contradictory snapshots while patients wait longer for therapy initiation. This article suggests a flexible plan that brings together different data systems with a focus on time-based modeling. Lifecycle signals come together using standard event-sourced structures, creating operational views that can be accessed through multiple channels at the same time. Saga coordination patterns merge with command-query separation strategies for multi-step process management. The system protects privacy by using access methods based on user attributes and roles, which work together with rules for handling data and secure logging. Transaction capacity meets enterprise demands while propagation speeds approach real-time thresholds. An accompanying measurement approach links architectural choices with service quality indicators. Observational data from deployment scenarios revealed shortened call durations for status inquiries and fewer repeated contacts. Such evaluation methods permit rigorous assessment of machine-learning-enhanced agent tools without sacrificing regulatory adherence.

1. Introduction

1.1 Operational Landscape and Driving Forces

Patient service quality in pharmacy and payer settings cannot be separated from clinical transaction flows and financial processing streams. A simple question from a patient about whether their medication is approved starts a series of steps that involve checking authorization records, benefit policy databases, adjudication logs, shipping lists, assistance program eligibility, and contact preferences. These information pools typically occupy isolated platforms. Service agents consequently toggle between applications, assemble fragmentary histories, and untangle conflicting data points under tight time constraints. Operational drag emerges through delayed medication access, inflated service delivery expenses, and amplified compliance vulnerabilities. Research spanning 252,145 patient encounters across 285 healthcare

facilities demonstrated that service center performance directly shapes patient perceptions of access quality and satisfaction levels [7].

Prescription tracking systems that focus on events have shown that organizing medication processes in a clear timeline helps measure everything more effectively. Therapy start timelines can be tracked, delays can be linked to specific steps in the process, and it becomes possible to use predictive algorithms to address issues like authorization rejections. Organizations can closely monitor prescription populations using this event-centric model, eliminating the need for delicate manual matching exercises. This makes operations more open, from the outside into the architectural center. Healthcare information exchange specifications establish foundational vocabularies for assembling unified perspectives across disparate clinical platforms [1]. Federal directives targeting system connectivity and authorization streamlining simultaneously generate compliance obligations

and transformation opportunities for pharmacy service infrastructures [2].

1.2 Prior Research and Manuscript Contributions

Earlier investigations addressing healthcare platform architecture and data infrastructure emphasized capability taxonomies, medical record connectivity, and reporting repositories, generally prioritizing clinical datasets over instantaneous service center demands. Customer information platforms and relationship management tools similarly sharpen visibility for commercial engagements yet typically bypass event-sourced underpinnings, graph-oriented temporal structures, and saga-driven coordination suited for regulated pharmacy processes. As a result, prevailing frameworks underserved unified, low-latency prescription tracking, service-center-tuned projections, and consent-governed access within privacy-regulated environments that must comply with connectivity mandates.

This manuscript makes four contributions:

1. It specifies a unified data foundation blueprint that combines event-sourcing, CQRS, and saga coordination tailored to omnichannel pharmacy and payer service centers.
2. It operationalizes graph-based temporal structures as bounded subdomains for risk, fraud, and care coordination while keeping high-volume status retrieval on tabular projections.
3. It links architectural mechanisms to observable service-center indicators (throughput, currency, agent productivity, reliability) using a before–after deployment assessment.
4. It outlines a phased deployment and governance sequence aligning technical capabilities with data stewardship and regulatory constraints in multi-party ecosystems.

The blueprint reflects extensive real-world experience in payer and specialty pharmacy settings where service platform governance has been exercised under large-scale transaction volumes.

2. Operational Challenges and Target Outcomes

2.1 Characterizing Unified Healthcare Data Foundations

A unified healthcare data foundation acts like a valuable resource, collecting information from various sources to create clear and organized views

of current operations and past trends. Unlike traditional reporting warehouses that focus on processing data in large batches, these foundations are designed for quick and efficient data access to support agents, automated systems, member websites, and coordination tools. The architectural method described here comes from working with multiple programs across national insurance companies and specialty pharmacies, where service centers deal with a large number of prescription and benefit questions while following strict regulations. Healthcare organizations handling 2.25 million patient visits annually alongside 3.4 million distinct medical transactions have demonstrated that unified data foundations substantially reduce operational complexity while improving service delivery consistency [8]. Standardized prescription messaging specifications play indispensable roles enabling cross-platform consistency within hub designs [3].

2.2 Target Outcome Specifications

Healthcare technology leadership, pharmacy benefit executives, and service center administrators consistently voice five outcome targets for data foundation initiatives. Capacity targets state that systems must handle a large number of daily transactions and be able to manage sudden increases in demand caused by new treatments, policy changes, or seasonal changes. Currency targets mean that important updates need to reach staff and online services almost immediately, reducing unnecessary work and follow-up calls. Studies looking into how easy it is to get to healthcare services by phone found that facilities that cut their average response times from 87 seconds to 69 seconds also saw their call abandonment rates drop from 12.0% to 8.3%. Almost 78% of patients said they were able to get urgent care at high-performing centers [7]. Reliability targets mandate consistent state portrayal across telephone, chat, messaging, and self-service touchpoints, preventing channel inconsistencies that frustrate patients and inflate contact volumes.

Stewardship targets ensure privacy by limiting access to only what is necessary, using consent for data handling, and having audit features that fit Technical protection and activity logging mandates under healthcare privacy statutes shape architectural constraints demanding satisfaction [4]. Coordination targets require orchestration of multi-phase processes spanning benefit verification, authorization handling, claim filing, and exception resolution with tolerance for partial breakdowns and visibility into progression status.

3. Architectural Blueprint

3.1 Structural Layering and Component Arrangement

The proposed blueprint combines an event-focused data foundation with a coordination layer that executes business workflows under strong monitoring and audit controls. The foundation ingests signals from prescribing networks, benefit determination engines, authorization platforms, claims processors, dispensing operations, and engagement channels, normalizes them into standard formats, enriches them with reference data, and persists them as secure event histories. These histories feed read-optimized projections and time-aware structures that omnichannel consumers access through secure programming interfaces, while coordination components subscribe to event streams to trigger follow-up actions. Prior work on command–query separation in medical information systems demonstrated that reducing schema complexity from 17 tables with 178 columns to 5 tables with 54 columns can improve response times by nearly 40% and cut data transfer sizes to about one-third, illustrating the performance potential of similar design choices [8].

Structural layering partitions concern absorption, enrichment, storage, projection, and consumption strata. Security laws provide rules for controlling data flow and access at each layer. This separation allows each part to scale independently, receive targeted optimization, and maintain clear ownership, while preserving the complete audit history required for regulatory compliance.

3.2 Canonical Structures and Connectivity Standards

A unified data foundation must harmonize heterogeneous ecosystems spanning version-two clinical messaging from treatment platforms, electronic interchange transactions for eligibility and claims handling, vendor-particular authorization sequences, and service center engagement signals. A workable reference strategy establishes canonical structures for shared meanings while letting domain groups publish data assets evolving autonomously. Connectivity standards have achieved broad uptake as canonical formats for patient, encounter, medication, and financial objects, serving as core semantic strata for patient-therapy-benefit-outcome perspectives.

Translating legacy message structures and interchange segments into standard-aligned assets enables uniform projection schemas and smooths

downstream programming interface delivery for member access and provider sequences. Where connectivity standards prove unsuitable, such as routing signals and telephony markers, the foundation can embrace complementary schemas carrying robust identifiers and stewardship metadata. Foundational investigations examining service center operations, queue behaviors, Throughput optimization provide operational backdrops within which canonical structures must operate [6].

3.3 Event-Oriented Absorption and Projection Generation

Event-oriented pipelines shrink polling delays and multi-platform query sprawl burdening agent applications. Instead of checking many platforms for each interaction, the system calculates the current status in advance and provides it through projections that are customized for specific use cases. Core absorption tasks involve getting rid of duplicates, handling repeated actions, organizing data for timely use, merging information from key datasets and policy sources, and keeping track of data history from source platform tags, absorption times, enrichment versions, and projection updates. Projections need to be designed to meet the needs of service centers and coordination, including summaries of current status, important timelines, and displays for exceptions and tasks that are waiting. The projection layer acts as the main point of contact for agent workstations and chat tools, delivering responses in less than a second for common queries while staying in sync with the main event databases.

3.4 Graph-Based Temporal Structures

Healthcare service quality requires understanding the connections between different factors like bonds, therapy changes, and signs of fraud or risks of escalation. A single data source can start as a collection of customer information that gathers signals from various channels into a timeline of patient views. Smart changes to graphs can reveal important connections, like which doctors and pharmacies work together or how different conditions relate to medications, which would be too costly to determine using regular data tables.

Graph repositories can introduce notably elevated operational burdens regarding storage footprints, replication demands, and query intricacy. A reference strategy accordingly reserves graphs for bounded subdomains spanning risk and fraud detection alongside care coordination while preserving primary transaction standing and

chronologies in event repositories paired with tabular projections. This selective stance balances analytical strength from graph-oriented relationship navigation against operational straightforwardness and cost containment of tabular projections for high-traffic, low-complexity retrievals.

4. Protection, Privacy, and Stewardship

4.1 Activity Logging and Access Governance

Healthcare data foundations represent appealing targets given their concentration of sensitive personal and financial records. Architecture must weave protection controls into data pathways rather than appending them subsequently. Statutes mandate mechanisms for capturing and reviewing platform activity touching electronic protected health records. Event-oriented architectures naturally meet these requirements when the signals are secure, the history is preserved, and the access activities are properly monitored. A study that looked at 252,145 patient visits in 285 healthcare facilities found that how quickly service centers answer phone calls is closely linked to how patients feel about the quality of access, with facilities that had slower average response times showing higher satisfaction scores for urgent care access.

A workable access governance pattern combines role-oriented permissions that differentiate agents, pharmacists, and analysts with attribute-oriented rules that cover minimum necessary retrieval, consent boundaries, and residency markers. Analytical services can obtain pseudonymized markers, while agent workstations obtain patient contact particulars needed for outreach. This stratified stance lets the foundation serve varied consumer groups with fitting data exposure while preserving unified underlying data architectures.

4.2 Consent Handling and Distributed Stewardship

Consent markers demand treatment as primary signals shaping projection assembly. Consent-aware projections help users avoid accidentally combining restricted attribute sets, making sure privacy is part of the system's design instead of just depending on rules at the application level. This approach ensures that any changes to consent automatically apply to all related projections without needing manual updates or allowing temporary inconsistencies.

Multi-party ecosystems distribute data custody across payers, pharmacies, providers, and patients. A distributed data mesh approach allows these groups to keep control of their data and service

responsibilities, while the shared system acts as a manager who ensures data history, agreement on data structure, and rule enforcement. Using a command-query separation pattern has improved the speed of medical information platforms while still keeping the necessary controls for following regulations.

5. Assessment and Deployment

5.1 Assessment Framework and Indicator Selection

The proposed data framework was assessed in large enterprise service centers handling status and eligibility traffic across prescription and benefit sequences using a quasi-experimental, before–after design. Baseline intervals preceding foundation activation were compared with post-deployment intervals, with weekly matching by therapy composition, payer composition, and channel split across telephone and digital interactions. Status-exclusive contacts encompassed interactions where no downstream sequence beyond information delivery received initiation, and repeated contacts captured any subsequent engagement for identical cases within specified windows.

Assessment concentrated on four indicator groupings. Throughput and currency captured the duration from source signal emission to projection refresh under realistic peak traffic and breakdown scenarios. Agent productivity covered average handle time for status and eligibility calls, repeated-contact frequency for the same cases, and the success rate of resolving status inquiries in a single contact. Reliability and correctness were evaluated by comparing projections to underlying source states in the presence of duplicates, incorrect signals, and retries, using markers such as deduplication tags. Operational overhead was monitored through escalation frequency and the accumulation of manual tasks for approvals and exception handling, providing a view into residual friction after coordination automation. Related investigations of healthcare service centers, which reported reductions in average response times from 87 seconds to 69 seconds and call abandonment decreases from 12.0% to 8.3%, illustrate how changes in operational metrics can influence patient perceptions of access and satisfaction [7].

5.2 Deployment Observations and Interpretations

Deployments of the enterprise service center that provided status and benefit information through unified foundation queries significantly reduced the

average time taken to handle status and eligibility inquiries during busy periods. Repeated-contact frequencies within specified windows were contracted observably across identical intervals. These changes happened when agents stopped switching between different platforms and instead used a single, organized system, while the number of staff and how calls were directed remained the same. Observations demonstrate how architectural shifts at data and coordination strata convert directly into quantifiable service quality and operational gains. Similar results were seen in medical information systems where separating commands and queries cut the time for admission forms from 1.243 seconds to 0.774 seconds and also lowered the amount of data retrieved from 32.296 kilobytes to 11.658 kilobytes for each transaction.[8]

Interpretation of these observations requires recognition of several constraints. The deployment examples reflect specific conditions in large enterprise service centers and may not generalize to all payer, pharmacy, or integrated delivery settings. Although primary staffing counts and routing configurations were held constant during the observation windows, unmeasured operational shifts may have influenced outcomes. The analysis emphasizes near-term operational indicators and does not directly capture longer-term clinical, financial, or patient-reported satisfaction effects beyond what is inferred from related literature. Event-sourcing patterns nonetheless provide a consistent framework for constructing secure, auditable event records that support such observational assessments.

5.3 Deployment Sequencing

Transformation toward unified data foundation capabilities constitutes a socio-technical undertaking demanding synchronized technical and organizational shifts. A pragmatic sequencing typically advances through four phases. The first phase establishes visibility by collecting core signals about authorization decisions, adjudication outcomes, and shipping details, and exposing consolidated timelines to agents under strict access

controls. The second phase broadens canonical reach by extending mappings to relevant connectivity standards, layering master-data enrichment, and regularizing signal vocabularies across domains.

The third phase introduces coordination using saga-oriented sequences to check benefits, route authorizations, and handle exceptions, emphasizing rules that are transparent and maintainable. The fourth phase selectively augments graph capabilities and machine-learning support by applying bounded graph subdomains to risk detection, fraud identification, and care coordination, while embedding monitoring and human oversight for ML-assisted decisions. Customer information platform concepts inform the unified foundation strategy for consolidating omnichannel signals [11].

5.4 Organizational Transition Planning

Organizations advancing through at least the opening three phases have exhibited consistent tendencies. Visibility-focused baselines lead to quicker reductions in repeated status-only communications by providing agents with complete information through combined displays. Canonical broadening enables more dependable cross-channel chronologies by guaranteeing consistent meanings across telephone, digital, and self-service touchpoints. Saga-oriented coordination shrinks manual rework for authorization and exception handling by automating routine progression and highlighting exceptions for human judgment. Organizational change requires clear planning because foundational programs often succeed when they align the goals of different departments, including operations, compliance, technology, and clinical teams. Assigning custody for data assets and schemas fixes accountability for data integrity and evolution. Creating groups to oversee schema updates and policy changes ensures that the foundation can quickly respond to business needs while maintaining the necessary consistency and compliance required in regulated healthcare environments.

Table 1: Healthcare Service Center Operational Challenges [7]

Challenge Area	Description
First Point of Contact	Telephone often serves as initial patient engagement with healthcare system
Service Scope	Assists with appointments, medical questions, billing, administrative tasks
Management Model	Overwhelming majority (93%) managed in-house
Growth Expectation	Service levels, staffing, and telephone-based service importance predicted to increase
Regulatory Requirements	Federal health insurance exchanges required to operate telephone hotlines

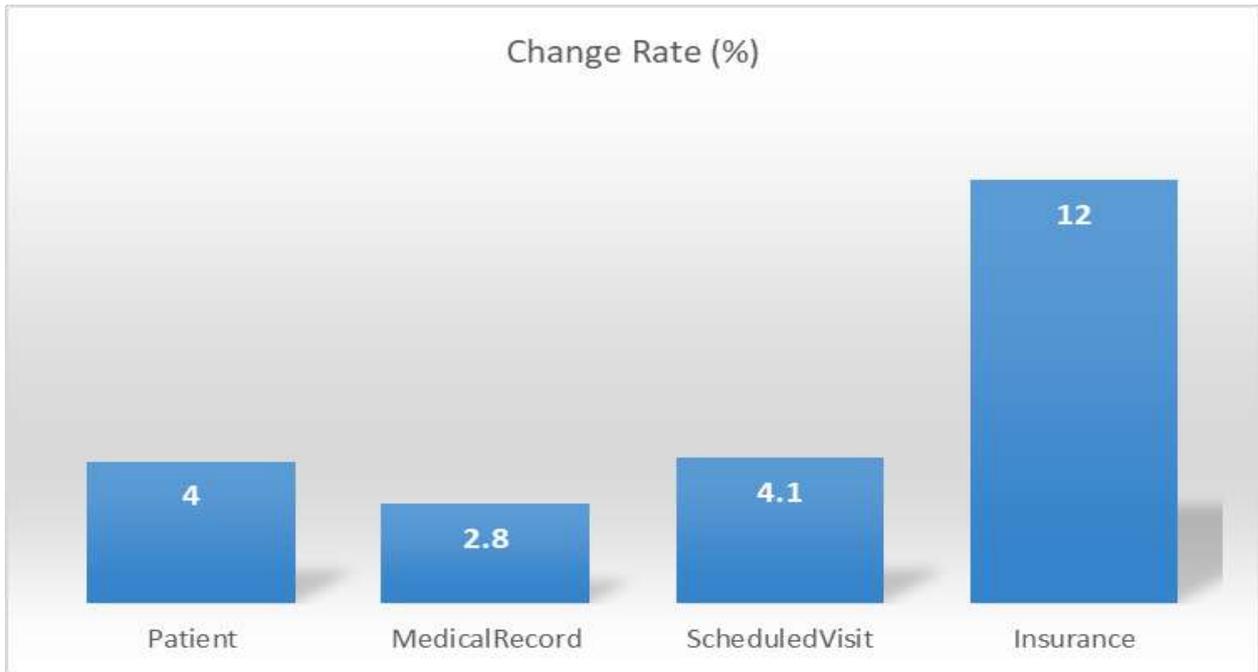


Figure 1: Data Change Frequency in Healthcare Information Systems [8]

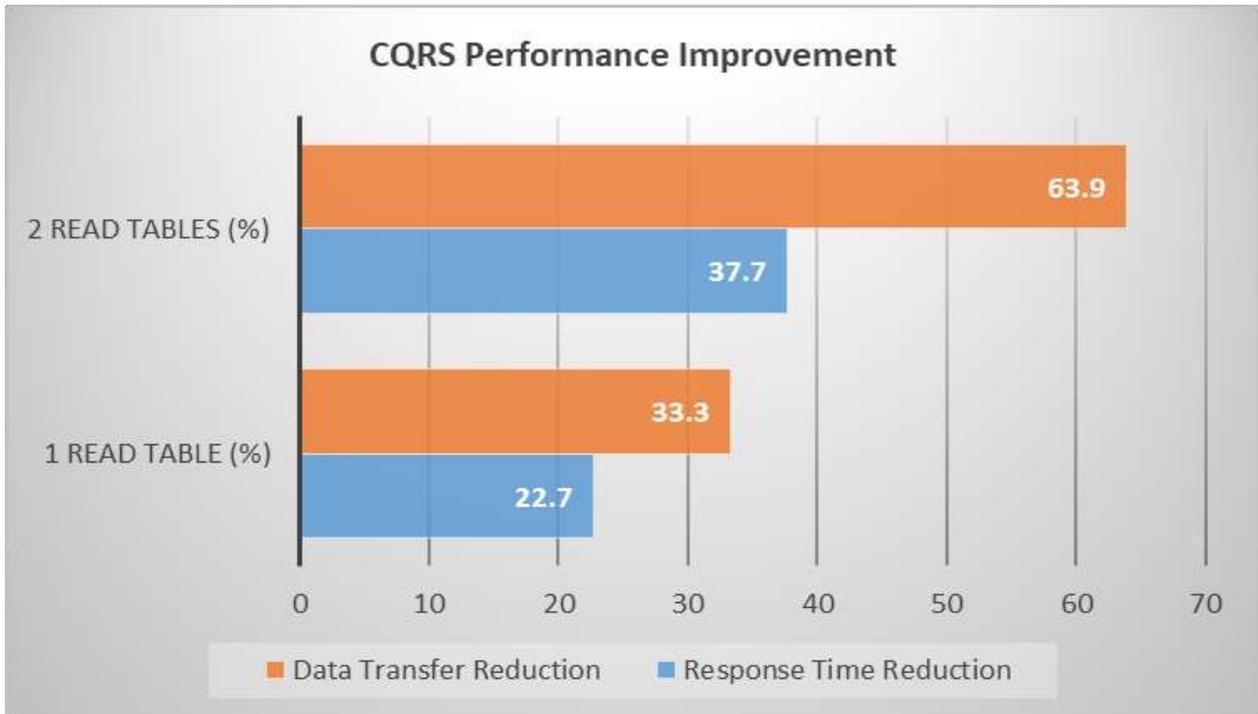


Figure 2: CQRS Performance Improvement [8]

Table 2: Command-Query Responsibility Segregation Architectural Components [9]

Component	Function	Characteristics
Command Side	Processes insert, update, delete operations	Transaction-focused, normalized data storage (3NF), consistent data handling
Query Side	Handles read/select operations	Read-optimized, denormalized data storage (1NF), eventually consistent
Thin Read Layer	Projects data directly to DTOs	Bypasses domain model, connects directly to data storage, no impedance mismatch
Synchronization Component	Maintains consistency between data stores	Event-driven updates handle data transformation between models

6. Conclusions

Unified healthcare data foundations paired with graph-based temporal structures present a sound architectural answer to the core service center predicament of scattered standing across platforms functioning acceptably alone yet yielding collective confusion. Embracing event-sourced underpinnings, canonical modeling, tracking connectivity standards, query-tuned projections, and saga-oriented coordination lets enterprises furnish consistent omnichannel encounters while reinforcing audit capabilities and privacy governance fitting regulated healthcare settings.

The plan ensures that different vendors can work together and focuses on measuring results, so improvements like faster access, fewer repeated contacts, shorter handling times, and safer machine-learning upgrades can be confirmed instead of just claimed across payer, pharmacy benefit, specialty pharmacy, and integrated delivery areas. Creating a common data system and using graph-based time structures as adaptable models for new connectivity and authorization rules provides a guide for transforming service centers to be ready for machine learning.

The architectural blueprint exhibits how the design of an enterprise healthcare platform can merge service center tools, coordination engines, and data stewardship into coherent, measurable capabilities boosting visibility, shrinking friction, and backing safe rollout of machine-learning-aided sequences. The unified foundation serves as a consolidating source that decouples workstations and conversational tools from scattered transaction repositories within enterprise service platforms. The blueprint operationalizes Command-query separation and saga patterns allowing agents and supervisors to directly observe service centers for coordination purposes, witnessing not only the current standing but also active sequence phases, executed fallback routes, and gathering exceptions. Future studies could expand this evaluation system to compare multiple locations, use random or semi-experimental setups, and track longer-term clinical, financial, and satisfaction outcomes. Healthcare organizations persist in facing demands for upgraded patient access, trimmed operational outlays, and strengthened regulatory conformance. The unified healthcare data foundation blueprint furnishes a principled groundwork for enduring transformation balancing novelty with stewardship controls indispensable for sustaining confidence in healthcare information platforms.

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.
- **Use of AI Tools:** The author(s) declare that no generative AI or AI-assisted technologies were used in the writing process of this manuscript.

References

- [1] HL7, "FHIR R5 (v5.0.0) Specification," 2025. Available: <https://hl7.org/fhir/R5/>
- [2] Centers for Medicare & Medicaid Services, "CMS Interoperability and Prior Authorization Final Rule (CMS-0057-F)," 2025. Available: <https://www.cms.gov/cms-interoperability-and-prior-authorization-final-rule-cms-0057-f>
- [3] National Council for Prescription Drug Programs (NCPDP), "HIPAA Information: NCPDP Telecommunication Standard Version F6 and compliance timeline," 2025. Available: <https://ncpdp.org/NCPDP/media/pdf/Resources/HIPAA.pdf>
- [4] Legal Information Institute, Cornell Law School, "45 CFR §164.312 Technical Safeguards (Audit Controls)," current. Available: <https://www.law.cornell.edu/cfr/text/45/164.312>
- [5] U.S. Department of Health & Human Services, "HIPAA Security Series #4: Technical Safeguards," 2007. Available: <https://www.hhs.gov/sites/default/files/ocr/privacy/hipaa/administrative/securityrule/techsafeguards.pdf>
- [6] N. Gans, et al., "Telephone Call Centers: Tutorial, Review, and Research Prospects," 2003. Available: <https://pubsonline.informs.org/doi/10.1287/msom.5.2.79.16071>
- [7] K. N. Griffith, et al., "Call Center Performance Affects Patient Perceptions of Access and Satisfaction," 2019. Available: <https://pmc.ncbi.nlm.nih.gov/articles/PMC8177735/>
- [8] Petar Rajković, et al., "Using CQRS Pattern for Improving Performances in Medical Information

- Systems," 2013. Available: <https://ceur-ws.org/Vol-1036/p86-Rajkovic.pdf>
- [9] G. Young, "CQRS Documents," 2010. Available: https://cqrs.files.wordpress.com/2010/11/cqrs_documents.pdf
- [10] Microsoft, "Event Sourcing pattern," Azure Architecture Center, 2025. Available: <https://learn.microsoft.com/en-us/azure/architecture/patterns/event-sourcing>
- [11] Scott Clark, "What Is a Customer Data Platform (CDP)? 2025 Market Insights," CMSWire, 2025. Available: <https://www.cmswire.com/customer-data-platforms/what-is-a-customer-data-platform-cdp/>