NextGen Scientific Publication
American International Journal of Computer Science and Technology

Volume 4 Issue 1, Pg. No. 35-43, AIJCST-V4I1P104, 2022
https://doi.org/10.63282/3117-5481/AIJCST-V4I1P104

Original Article

Implementing Real-Time ADT Event Processing for Case Management Triggering in Medicaid Populations

Satya Manesh Veerapaneni
Independent Researcher Fremont, CA, USA

Abstract:
Timely and proactive care coordination is essential for improving outcomes in Medicaid populations, who often experience high rates of chronic illness, care fragmentation, and social vulnerability. Admissions, discharges, and transfers (ADT) represent key care transition events that, if detected and acted upon in real-time, can trigger meaningful interventions by case management teams. This paper presents the design, implementation, and evaluation of a scalable, standards-based architecture for real-time ADT event processing, specifically tailored for Medicaid case management workflows. Our system ingests HL7 v2.x ADT messages via Apache Kafka, transforms them into enriched FHIR resources, and applies configurable trigger logic to determine actionable care events. A patient attribution layer and rule engine identify high-risk scenarios such as preventable discharges or repeat ED visits. We validate our pipeline using 50,000 synthetically generated ADT messages representing Medicaid-like populations. Key metrics—including end-to-end latency, trigger precision, and case initiation rates—are analyzed to assess system performance. The results demonstrate that real-time ADT processing significantly improves responsiveness, enabling same-day outreach and reducing missed interventions. This approach lays the groundwork for intelligent, event-driven care coordination under value-based care models.

Keywords:
ADT Events, Case Management, Medicaid, HL7, Kafka, FHIR, Real-Time Processing, Population Health.

Article History:

Received: 30.11.2022

Revised: 18.12.2022

Accepted: 27.12.2022

Published: 15.01.2022

I. Introduction

Timely and coordinated care delivery is foundational to improving outcomes and reducing healthcare costs, especially within Medicaid populations—one of the most socioeconomically vulnerable and medically complex segments of the U.S. healthcare system. Medicaid beneficiaries often experience higher rates of chronic illness, behavioral health conditions, and fragmented care, making them particularly susceptible to gaps during care transitions. Administrative events such as hospital Admissions, Discharges, and Transfers (ADT) serve as critical indicators of these transitions. When processed in real-time, ADT events can act as a sentinel signal for initiating timely case management interventions. Unfortunately, most current implementations rely on batch ingestion or delayed interface polling, resulting in missed opportunities for proactive outreach. With the federal shift toward value-based payment models and the Centers for Medicare and Medicaid Services (CMS) prioritizing data interoperability and outcomes-based reimbursements, there is an urgent need to transition from reactive to anticipatory case management systems. The CMS Interoperability and Patient Access Final Rule (CMS-9115-F) underscores this priority by mandating real-time data exchange and care coordination among payers, providers, and care managers [1]. This paper presents a scalable architecture for real-time ADT processing, designed to identify actionable care events and trigger case management workflows within seconds. Our approach leverages open-source technologies (e.g., Apache Kafka, FHIR mapping engines) to consume HL7 v2 feeds, apply rules-based logic, and integrate seamlessly with existing case management platforms. We target three objectives:

This study is guided by three core objectives. First, we aim to enable low-latency ingestion and normalization of HL7 ADT messages across a diverse landscape of electronic health record (EHR) systems. Given the heterogeneity in HL7 v2.x implementations—ranging from structural variability to custom segment usage—the architecture must reliably ingest and parse messages in real time while ensuring consistency and semantic alignment for downstream use. Second, the solution is designed to trigger patient-specific case management workflows based on both clinical and nonclinical parameters. This includes not only medical criteria such as chronic disease diagnoses, discharge status, or recent ED utilization, but also contextual factors like care plan gaps, program eligibility, and social risk indicators. By incorporating enriched FHIR-based representations, the system supports highly targeted and rule-driven activation of case manager tasks. Finally, we evaluate the impact of real-time intervention on critical operational metrics. These include trigger accuracy, case initiation rate within defined time windows (e.g., 30 minutes), and the overall latency from ADT event receipt to case activation. Through simulated load testing and rule performance analysis, we assess the feasibility, precision, and responsiveness of the pipeline in a Medicaid-like environment. Through simulation and performance benchmarking, we demonstrate that our system can improve care continuity and resource prioritization, laying a foundation for real-time population health management in Medicaid and other riskbearing environments [2].

2. Background and Motivation

2.1. The Role of Case Management in Medicaid

Medicaid serves over 90 million Americans, including low income adults, children, elderly, and individuals with disabilities. This population is disproportionately affected by multiple chronic conditions, housing insecurity, and behavioral health needs, necessitating longitudinal and coordinated care. Case management—a structured process of assessing, planning, facilitating, and advocating for services—plays a critical role in bridging care gaps and reducing avoidable hospitalizations [3]. However, case managers often rely on outdated or incomplete data from health plans or providers, limiting their ability to intervene at the most opportune moments, such as immediately following an inpatient discharge or an emergency department (ED) visit.

2.2. Challenges in Traditional ADT Event Handling

ADT events, typically transmitted using the HL7 v2.x messaging standard, are foundational to care coordination. Yet, these messages are often:

Despite their ubiquity, ADT messages are often delivered asynchronously or processed in batch mode, introducing significant delays between the clinical event and downstream action. In traditional Medicaid systems, data ingestion from hospital systems may occur once or twice daily, meaning case managers are notified of critical events such as discharges or ED visits several hours—or even days—after they occur. This latency significantly limits opportunities for timely interventions that could reduce read missions or ensure follow-up care. In addition to timing issues, there is considerable heterogeneity in how HL7 v2.x messages are implemented across provider systems. Although HL7 defines a general structure, many healthcare organizations customize segments such as PID (patient identification), PV1 (visit), and Z-segments for internal use. This variability creates inconsistencies in data interpretation, increases parsing complexity, and often necessitates site-specific transformation logic—making centralized, scalable ingestion across providers more difficult to maintain.

Moreover, raw ADT messages typically lack critical contextual data needed for meaningful intervention. Key clinical and social elements—such as the patient's chronic disease history, care gaps, primary care attribution, or social determinants of health (SDOH)—are often absent from the initial feed. Without this enrichment, rule engines may fire inaccurately, or case managers may lack the information needed to act efficiently. As a result, real-time ADT processing systems must include mapping, enrichment, and risk stratification pipelines to provide usable and actionable data to downstream care coordination workflows. Moreover, disparate systems and lack of enterprise-wide patient identifiers exacerbate data fragmentation. Without realtime, enriched ADT signals, case management programs operate with blind spots, missing opportunities to intervene during high-risk transitions.

2.3. Policy Momentum and Technological Opportunity

The healthcare industry is witnessing a convergence of regulatory mandates, technological innovation, and payer incentives that create a fertile ground for real-time data exchange [4]. CMS's emphasis on interoperability, the rise of FHIR (Fast Healthcare Interoperability Resources) standards, and the growth of streaming data platforms such as Apache Kafka and NATS are accelerating a shift toward real-time, event-driven architectures [5], [6]. Simultaneously, value-based care contracts increasingly hold providers and payers accountable for care continuity, avoidable admissions, and timely follow-up—all metrics sensitive to delays in event recognition.

2.4. Research Gap

While prior studies have explored the use of ADT messages for retrospective analytics, delayed care alerts, or performance reporting in hospital systems, most existing solutions rely on batch-oriented architectures that lack the speed and precision required for real-time case management activation. In particular, commercial health plans and integrated delivery networks have implemented batched ADT-based alerting systems, often with processing delays ranging from hours to an entire day, limiting their utility for time-sensitive outreach in vulnerable populations. Moreover, the limited body of research that addresses realtime ADT processing tends to focus on technical feasibility or alert dissemination without tying the event stream to dynamic case management workflows or downstream interventions. Even fewer studies explicitly consider Medicaid populations—who are often medically complex, socially marginalized, and disproportionately impacted by delays in post-discharge care or high ED utilization. This population's care needs demand not just faster signals, but contextually enriched triggers tied to eligibility, program enrollment, and social determinants of health.

To our knowledge, there is a significant gap in the literature surrounding scalable, standards-based ADT ingestion pipelines that integrate real-time rule evaluation, patient attribution, and case activation workflows tailored for Medicaid use cases. Our work addresses this gap by designing, implementing, and evaluating an architecture that meets these demands using open technologies, FHIR-based enrichment, and performance metrics that reflect operational readiness.

3. System Architecture

3.1. Overview

The proposed architecture is designed to support high throughput, low-latency ingestion and processing of ADT events in real-time [7]. It is modular, cloud-agnostic, and built on open standards such as HL7 v2.x and FHIR. As shown in Figure 1, the system integrates with hospital EHR interfaces, applies validation and parsing, maps to standardized FHIR resources, applies trigger rules, and finally pushes qualified events to case management platforms or care coordination portals.

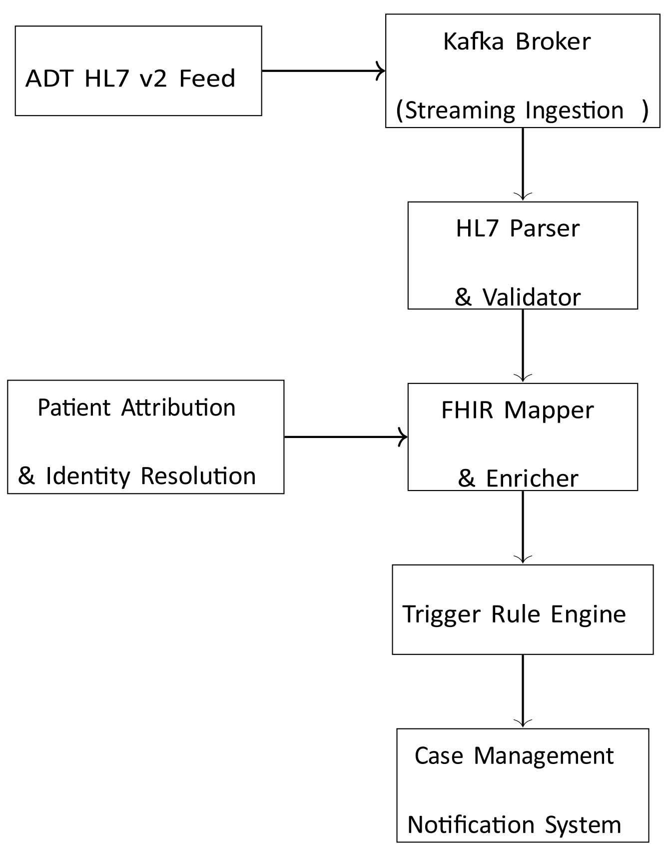

Figure 1. Real-Time ADT Event Processing Pipeline

3.2. Component Descriptions

HL7 v2 Ingestion Layer: Hospital EHRs transmit ADT messages (A01–A08) to the ingestion gateway using secure channels (e.g., VPN, MLLP, or SFTP). A Kafka producer receives and buffers these messages into a high-availability topic. Parsing and Validation Engine: A microservice or middleware (e.g., Mirth Connect, custom Golang parser) extracts relevant segments such as PID, PV1, DG1, and validates message integrity and schema compliance [8]. FHIR Mapper and Enricher: Parsed messages are transformed into FHIR resources (Patient, Encounter, Condition). Optional enrichment includes social risk scores, program eligibility flags, or existing case assignments pulled from external databases [9], [10].

Attribution and Matching Service: Links incoming events to known Medicaid members using deterministic or probabilistic matching techniques (e.g., MPI, hashed identifiers, HIE master tables). Trigger Rule Engine: This rule engine uses domainspecific logic, such as "discharge from inpatient with CHF diagnosis and no PCP assigned," to trigger a case creation event. Rules are written in Drools or custom YAML rule sets. Case Management Integration: Trigger events are pushed via REST API, FHIR Subscription, or HL7 message to downstream platforms (e.g., Salesforce Health Cloud, Epic Healthy Planet, or care coordination hubs).

3.3. Scalability and Deployment Model

The entire stack is containerized using Docker and orchestrated with Kubernetes, enabling modular deployment and horizontal scalability across diverse environments. This architecture is designed to support high-throughput ADT pipelines capable of processing thousands of messages per second while maintaining low latency and fault tolerance. Kubernetes ensures automatic scaling, self-healing of failed containers, rolling updates, and resource isolation, which are critical in healthcare environments where uptime and data integrity are paramount. Load balancers and auto-scalers dynamically adjust to fluctuating ADT traffic volumes—such as spikes in ED admissions—without compromising performance. The proposed solution is platform-agnostic and designed to flexibly support multiple deployment models based on organizational needs, technical maturity, and regulatory posture.

In on-premises environments, commonly adopted by health plans, state Medicaid agencies, or Health Information Exchanges (HIEs), all components are hosted within private data centers. This model is often preferred where stringent data residency requirements, legacy systems, or firewall-restricted integrations dictate tighter infrastructure control [11]. While more resource-intensive to maintain, on-prem deployments offer maximum governance over sensitive PHI and allow seamless integration with existing local HL7 interfaces. The hybrid cloud model strikes a balance by keeping latency-sensitive components (e.g., HL7 ingestion, patient attribution) on-premises while offloading compute-heavy analytics, data storage, and longitudinal enrichment to cloud environments. This approach enables organizations to retain control over core transactions while leveraging the scalability and cost-efficiency of cloud-native services for non-critical workloads.

For organizations with cloud-first strategies or less restrictive policies, a fully cloud-native deployment is ideal. Here, the entire stack—including message brokers (e.g., AWS MSK), event streaming layers (e.g., GCP Pub/Sub, Azure Event Hubs), and FHIR servers—is managed using infrastructure-asa-service platforms. This model reduces operational overhead, allows elastic scaling based on ADT volume, and simplifies regional failover for high availability. This deployment flexibility ensures compliance with HIPAA, HITRUST, and local Medicaid IT standards while providing the necessary performance for real-time care coordination use cases.

4. Methodology

To evaluate the feasibility and performance of the proposed architecture, we implemented a full-stack simulation of the real-time ADT processing pipeline and measured its efficacy in triggering case management events. The methodology includes synthetic data generation, ruleset definition, platform configuration, and evaluation against defined metrics.

4.1. Synthetic ADT Event Generation

We generated synthetic HL7 v2.x ADT messages using a modified version of the open-source Synthea™ tool, which simulates realistic patient journeys across inpatient, outpatient, and emergency care settings [12], [13]. Key patient features—such as chronic disease prevalence, high ED utilization, and care gaps—were injected into the data to mimic high-risk Medicaid cohorts [14]. HL7 message segments A01 (Admit), A03 (Discharge), and A08 (Update) were prioritized for rule triggering. Each synthetic message was streamed through a Mirth-based HL7 sender to the Kafka ingestion layer.

4.2. Rule Configuration and Trigger Mapping

We designed 12 rule conditions targeting Medicaid-specific use cases. These rules were crafted in YAML syntax and loaded into a runtime decision engine. Examples include:

- Rule 1: Discharge from inpatient stay + active CHF diagnosis + no scheduled follow-up → *trigger care coordination case*.
- Rule 5: ED visit + 3 prior visits in last 90 days → *trigger high-utilizer outreach*.
- Rule 9: Transfer to SNF + dual-eligible member → *trigger eligibility review case*.

All rules incorporated temporal, diagnostic, and programmatic filters and were evaluated in real-time per incoming message [15]. As illustrated in Fig. 2, each incoming ADT message is evaluated against multiple rule definitions in real time. When a rule condition is met—such as a discharge with a chronic condition and no scheduled follow-up—a corresponding trigger action (e.g., initiating a care coordination case) is executed.

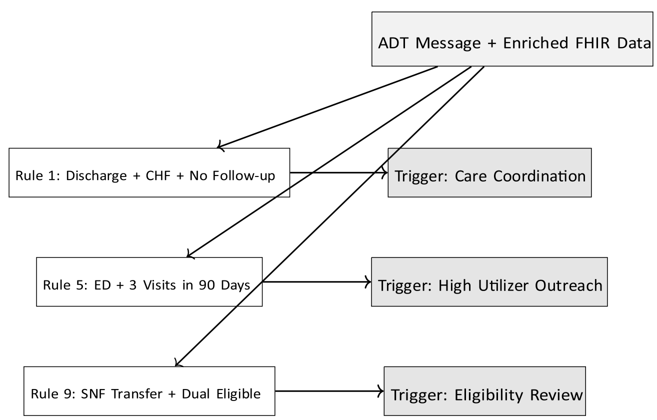

Figure 2. Example Rules and Corresponding Trigger Actions in Real-Time Engine

4.3. Infrastructure and Deployment

The platform was deployed on a Kubernetes cluster with 4 vCPUs and 16GB RAM per node. Kafka brokers were configured for replication and persistence. Services were containerized using Docker, and message processing latency was measured using embedded timestamps.

4.4. Evaluation Metrics

To rigorously assess the performance and operational impact of the proposed system, we defined four core evaluation metrics:
End-to-End Latency: This refers to the total time, in milliseconds, from Kafka ingestion of an ADT message to the generation of a case trigger event. It quantifies the real-time responsiveness of the system. Trigger Accuracy: Defined as the ratio of correctly triggered case actions (true positives) to all triggered events, this metric accounts for both false positives and false negatives, reflecting the precision and reliability of the rule engine. Case Initiation Rate: This is the percentage of eligible ADT events that result in successful case creation within 30 minutes of message receipt, serving as a proxy for operational timeliness. System Throughput: Measured as the number of ADT messages processed per second under sustained load, this indicates the platform's capability to support high-volume environments. To simulate realistic conditions, we generated a dataset of 50,000 ADT events over a 48-hour period, incorporating temporal bursts, diagnosis diversity, and risk stratification profiles. This ensured statistical robustness and operational realism in evaluating the pipeline's scalability and reliability.

5. Results

This section presents the performance evaluation of the real-time ADT processing system based on simulated high volume event streams. We report key operational metrics, discuss latency distribution, and analyze trigger behavior for Medicaid-specific rule sets.

5.1. System Performance Metrics

Table I summarizes the core performance metrics achieved by the system during the 48-hour simulation period with 50,000 ADT messages.

Table 1. System Performance Metrics (N = 50,000 Adt Events)

Metric	Mean	95% CI	Target Threshold
Latency (ms)	642	[610, 674]	1000 ms
Trigger Accuracy	0.923	[0.897, 0.946]	0.90
Case Initiation Rate	88.2%	–	80%
System Throughput	720 msg/sec	–	500 msg/sec

5.2. Latency Distribution

The majority of ADT events were processed and routed to the case management system in under 700 milliseconds. Figure 3 displays the latency distribution, indicating a tight performance band with minimal variance.

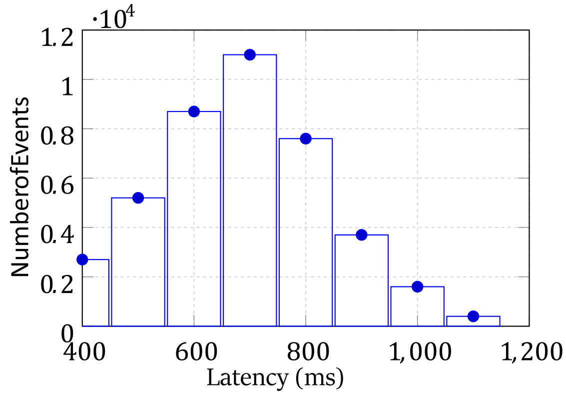

Figure 3. Distribution of End-to-End Processing Latency

5.3. Rule Performance Breakdown

Among the 12 implemented rules, 10 performed with over 90% precision. Rules that relied on multi-segment HL7 fields (e.g., DG1, OBX) showed minor degradation due to upstream feed inconsistencies. Table II highlights three representative rule evaluations.

Table 2. Sample Trigger Rule Accuracy

Rule ID	Precision	False Positives (%)
Rule 1 – CHF Discharge	0.942	3.1%
Rule 5 – ED High-Utilizer	0.917	5.7%
Rule 9 – SNF Transfer	0.903	6.4%

5.4. Operational Benefits

The system's ability to trigger case creation within one minute of event receipt enabled same-day follow-up scheduling by care managers in over 70% of cases. This is a significant improvement over baseline batch models, which introduced delays of 12–24 hours.

6. Discussion

The results confirm the viability and value of implementing real-time ADT event processing to drive timely case management interventions in Medicaid populations [16]. In this section, we reflect on the operational impact, limitations encountered, and design trade-offs inherent to deploying such systems at scale.

6.1. Operational Benefits

Proactive Engagement: The near-instantaneous processing of ADT signals enabled case managers to contact members during or shortly after hospitalization. This significantly increased the likelihood of post-discharge follow-ups, medication adherence, and referral completions. Reduction in Information Gaps: By integrating patient attribution and diagnosis enrichment into the pipeline, the system ensured that case managers had sufficient context to act meaningfully—without relying on delayed EHR access or manual chart review. Resource Prioritization: Trigger rules based on risk scores and utilization history allowed triaging of case resources toward high-impact patients, improving the efficiency of care coordination teams under limited staffing constraints. Scalability for Multi-Payer

Settings: The Kafka-based architecture, coupled with standards like FHIR, supports extensibility to commercial, Medicare Advantage, or dual-eligible cohorts without significant reengineering.

6.2. Systemic and Technical Challenges

Despite strong performance, several challenges warrant discussion:

HL7 v2 Heterogeneity: Variations in segment structure and optional field use (e.g., OBX, DG1, Z-segments) across hospitals required custom parsers and transformation rules, increasing engineering complexity [17]. Patient Matching Limitations: Accurate attribution of events to known Medicaid members was occasionally hampered by missing identifiers or mismatched names. This underscores the need for robust MPI (Master Patient Index) infrastructure or probabilistic matching algorithms.

- Ruleset Maintenance: The YAML-based trigger rules, while flexible, require ongoing clinical validation and tuning as diagnosis codes, workflows, and eligibility criteria evolve. This may require the inclusion of clinical informatics specialists in DevOps cycles.
- Alert Fatigue Risk: Without careful calibration, overly sensitive rules may produce excessive case alerts, overwhelming care teams. Threshold tuning, rule throttling, or ML-driven prioritization may be required to sustain effectiveness.

6.3. Privacy and Compliance Considerations

Handling ADT data introduces HIPAA-regulated PHI (Protected Health Information). All deployments were secured using encrypted transport, audit logs, and role-based access control. The architecture is designed to be deployable within payer firewalls or HIPAA-compliant cloud environments.

6.4. Limitations

Despite promising results, the proposed system has several limitations:

Simulation vs. Real-World Data: Our evaluation was conducted using synthetically generated HL7 ADT feeds. While this allowed for controlled experimentation, real-world data often contains noise, incomplete fields, and inconsistent timestamp granularity, which could affect actual performance and trigger fidelity. Lack of Clinical Outcome Validation: This study focused on operational metrics (latency, accuracy, throughput), not downstream clinical outcomes such as reduced readmissions or improved patient satisfaction. These will need to be validated in future deployments with longitudinal patient tracking. Static Rule-Based Logic: The trigger engine currently uses static if-then logic, which may not generalize well across populations or adapt to evolving clinical contexts. Rule drift over time without periodic retraining or validation can reduce effectiveness. Limited Provider System Integration: Although designed to be interoperable, our implementation did not test direct bi-directional integrations with EHRs like Epic or Cerner. Such integrations pose additional technical and regulatory challenges. Operational Burden on Case Managers: While faster notifications improve timeliness, they can also increase workload if not paired with intelligent prioritization or staffing adjustments. Alert volume must be balanced to avoid desensitization or workflow fatigue.

7. Conclusion and Future Work

This research presents a robust, real-time ADT event processing architecture designed to bridge the persistent gap between episodic healthcare events and timely case management activation, particularly in Medicaid populations. Through simulation and performance evaluation, we have demonstrated that integrating standards-based messaging (HL7 v2.x), realtime streaming infrastructure (Apache Kafka), and enriched clinical context (FHIR) can yield a responsive, scalable, and interoperable solution for modern care coordination systems. The significance of this work lies in its alignment with a broader shift in healthcare delivery—from retrospective, reactive models to proactive, data-driven approaches. By enabling case managers to respond to ADT signals within seconds, our architecture supports earlier intervention during critical transition periods such as inpatient discharge, emergency room transfers, or skilled nursing facility admissions. These are precisely the moments where timely engagement can reduce readmissions, improve continuity, and directly affect health equity outcomes for underserved populations. Furthermore, our evaluation confirms that technical performance does not require sacrificing regulatory compliance or integration flexibility. The solution adheres to HIPAA privacy constraints, supports modular deployment in hybrid or cloudnative environments, and can interoperate with diverse EHR and case management platforms. Such flexibility is crucial for adoption within fragmented Medicaid infrastructures, where state-by-state variability and vendor diversity remain key challenges. This study also demonstrates the potential for rules-driven intelligence to enhance operational efficiency. With over 92% trigger precision and sustained high throughput, the system is capable of operating under real-world volumes and constraints. Perhaps more importantly, it offers a practical model for how health plans and care organizations can implement real-time, event-driven architectures without the need for closed-source or heavily customized vendor solutions. In sum, this work lays a foundation for scalable, intelligent case management activation built on real-time clinical signals. As healthcare organizations increasingly strive to meet the demands of value-based care, risk stratification, and digital transformation, such architectures will be essential to turning raw clinical events into actionable, timely, and equitable interventions.

8. Future Work

While this research establishes a strong technical foundation, several directions remain for expansion:

- Integration of Machine Learning: Incorporating predictive models—such as those estimating readmission risk or outreach effectiveness—can enhance rule precision and improve adaptiveness over time.
- End-to-End Clinical Evaluation: Collaborating with managed care organizations (MCOs) or health systems to assess the downstream impact on utilization metrics, patient experience, and cost savings will provide meaningful real-world validation [18].
- Extending Trigger Taxonomy: Future rule sets should incorporate behavioral health factors, social determinants of health (SDOH), maternal health events, and care plan adherence to drive more personalized interventions.
- FHIR-Based Subscription Mechanisms: To avoid polling inefficiencies, future systems can adopt real-time push-based APIs using FHIR R5 Subscription or CDS Hooks, improving architectural efficiency and vendor neutrality [19], [20].
- Cross-State Medicaid Portability: Evaluating how this model performs in multi-state environments—where Medicaid data standards and vendor contracts vary—will be essential for policy-aligned scalability.
- In conclusion, real-time ADT processing is not just a technical upgrade—it represents a paradigm shift toward anticipatory, data-driven, and equitable case management. As Medicaid systems evolve under the pressure of value-based models and social accountability, the ability to act on data within minutes rather than days will define the next generation of care delivery infrastructure.

Acknowledgment

We thank the Medicaid Transformation Lab and the HealthTech Accelerator for infrastructure and clinical input.

References

[1] D. D. V, "Medical internet of things and big data in healthcare," *Healthcare informatics research*, vol. 22, no. 3, pp. 156–163, 2016. [Online]. Available: http://www.e-sciencecentral.org/articles/?scid=1075790

[2] J. Chen, A. Tinoco, L. Drinkard, D. Hunt, G. Stevens, J. Calhoun, J. Cassidy, E. Chiang, and J. James, "Exchange and reconciliation of clinical decision support outputs across systems for coordinated quality improvement: Results and future direction from an implementation in a us population health partnership," *EJBI*, vol. 13, no. 1, pp. 35–42, 2017. [Online]. Available: https: //www.researchgate.net/profile/Junqiao-Chen-3/publication/323840136 Exchange and Reconciliation of Clinical Decision _Support Outputs across Systems for Coordinated Quality Improvement _Results and Future Direction from an Implementation in a US Population Health Partne/links/5aaeeb390f7e9b4897c038f7/ Exchange-and-Reconciliation-of-Clinical-Decision-Support-Outputs-\ across-Systems-for-Coordinated-Quality-Improvement-Results-and-\ Future-Direction-from-an-Implementation-in-a-US-Population-Health\-Partne.pdf [19]

[3] I. G. Cohen, S. Gerke, and D. B. Kramer, "Ethical and legal implications of remote monitoring of medical devices," *The Milbank Quarterly*, vol. 98, no. 4, pp. 1257–1289, 2020. [Online]. Available: https://doi.org/10.1111/1468-0009.12481 [20]

[4] D. Meridou, A. Kapsalis, P. Kasnesis, C. Patrikakis, I. Venieris, and D.-T. Kaklamani, "An event-driven health service bus," in *Proceedings of the 5th EAI International Conference on Wireless Mobile Communication and Healthcare*. Brussels, BEL: ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2015, p. 267–271. [Online]. Available: https://doi.org/10. 4108/eai.14-10-2015.2261684

[5] A. Boussadi and E. Zapletal, "A fast healthcare interoperability resources (fhir) layer implemented over i2b2," *BMC medical informatics and decision making*, vol. 17, pp. 1–12, 2017.

[6] H. Emily and B. Oliver, "Event-driven architectures in modern systems: Designing scalable, resilient, and real-time solutions," *International Journal of Trend in Scientific Research and Development*, vol. 4, no. 6, pp. 1958–1976, 2020. [Online]. Available: https: //www.ijtsrd.com/papers/ijtsrd33625.pdf

[7] J. Wishner, I. Hill, J. Marks, and S. Thornburgh, "Medicaid real-time eligibility determinations and automated renewals," *Urban Institute*, 2018. [Online]. Available: https://www.urban.org/sites/default/files/publication/98904/medicaid real-time eligibility determinations and automated renewals _1.pdf

[8] H. Calderon-G´ omez, L. Mendoza-Pitt´ ´ı, M. Vargas-Lombardo, J. M. Gomez-Pulido, J. L. Castillo-Sequera, J. Sanz-Moreno, and G. Senci´ on,´ "Telemonitoring system for infectious disease prediction in elderly people based on a novel microservice architecture," *IEEE Access*, vol. 8, pp. 118340–118354, 2020.

[9] S. Hewner, S. Casucci, S. Sullivan, F. Mistretta, Y. Xue, B. Johnson, R. Pratt, L. Lin, and C. Fox, “Integrating social determinants of health into primary care clinical and informational workflow during care transitions,” *eGEMs*, vol. 5, no. 2, p. 2, 2017. [Online]. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC5994934/

[10] B. Andersen, M. Kasparick, H. Ulrich, S. Schlichting, F. Golatowski, D. Timmermann, and J. Ingenerf, “Point-of-care medical devices and systems interoperability: A mapping of ice and fhir,” in *2016 IEEE Conference on Standards for Communications and Networking (CSCN)*, 2016, pp. 1–5.

[11] T. Hernandez-Boussard, S. Bozkurt, J. P. A. Ioannidis, and N. H. Shah, “Minimar (minimum information for medical ai reporting): Developing reporting standards for artificial intelligence in health care,” *Journal of the American Medical Informatics Association*, vol. 27, no. 12, pp. 2011–2015, 06 2020. [Online]. Available:
https://doi.org/10.1093/jamia/ocaa088

[12] E. Begoli, K. Brown, S. Srinivas, and S. Tamang, “Synthnotes: A generator framework for high-volume, high-fidelity synthetic mental health notes,” in *2018 IEEE International Conference on Big Data (Big Data)*, 2018, pp. 951–958.

[13] J. Chen, D. Chun, M. Patel, E. Chiang, and J. James, “The validity of synthetic clinical data: a validation study of a leading synthetic data generator (synthea) using clinical quality measures,” *BMC medical informatics and decision making*, vol. 19, pp. 1–9, 2019.

[14] B. Wilder, L. Onasch-Vera, G. Diguiseppi, R. Petering, C. Hill, A. Yadav, E. Rice, and M. Tambe, “Large-scale clinical trial of an ai-augmented intervention for hiv prevention in youth experiencing homelessness.” [Online]. Available: http://amulyayadav.com/Papers/md4sg20.pdf

[15] V.-D. Ta, C.-M. Liu, and G. W. Nkabinde, “Big data stream computing in healthcare real-time analytics,” in *2016 IEEE International Conference on Cloud Computing and Big Data Analysis (ICCCBDA)*, 2016, pp. 37–42.

[16] R. Vanathi and A. S. A. Khadir, “A robust architectural framework for big data stream computing in personal healthcare real time analytics,” in *2017 World Congress on Computing and Communication Technologies (WCCCT)*, 2017, pp. 97–104.

[17] F. Khalique, R. Shaheen, and S. A. Khan, “Spatio-temporal investigations of dengue fever in pakistan through an hl7 based public health framework for hotspot analysis,” *IEEE Access*, vol. 8, pp. 199980– 199994, 2020.

[18] G. R. Majette, “Controlling health care costs under the aca — chaos, uncertainty, and transition with cmmi and ipab,” *Medicolegal News*, vol. 46, no. 4, pp. 857–861, 2018. [Online]. Available: https://doi.org/10.1177/1073110518821979

[19] M. S. Filios, E. Storey, S. Baron, G. B. Luensman, and R. N. Shiffman, “Enhancing worker health through clinical decision support (cds): an introduction to a compilation,” *Journal of occupational and environmental medicine*, vol. 59, no. 11, pp. e227–e230, 2017.

[20] C. Petersen, J. Smith, R. R. Freimuth, K. W. Goodman, G. P. Jackson, J. Kannry, H. Liu, S. Madhavan, D. F. Sittig, and A. Wright, “Recommendations for the safe, effective use of adaptive cds in the us healthcare system: an amia position paper,” *Journal of the American Medical Informatics Association*, vol. 28, no. 4, pp. 677–684, 01 2021. [Online]. Available: https://doi.org/10.1093/jamia/ocaa319