International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
ISSN:2147-6799 www.ijisae.org Original Research Paper

Adaptive RegTech Architectures for Real-Time Compliance in Modernized Financial Transaction Systems

Narasimha Rao Vanaparthi*[1], Maheshbabu Dhanekula*[2], Rishabh Srivastava*[3]

Abstract: Regulatory compliance in financial transaction systems has historically depended on batch-processing architectures that accumulate transaction records over defined intervals before applying rule-based inspection. Adequate for an era of infrequent regulatory examination and slower-moving financial crime, this model has become misaligned with a regulatory environment that now mandates continuous monitoring, near-real-time reporting, and immediate corrective action capability. This article proposes the Adaptive RegTech Architecture (ARA), a principled framework for achieving genuine real-time compliance in institutions operating hybrid infrastructure — specifically, organizations whose transaction processing cores remain mainframe-resident while compliance obligations require modern event-driven capabilities. The ARA is organized around four design principles: compliance layer decoupling via event-driven interfaces, event primacy as the authoritative source for compliance evaluation, unified graph-based entity identity management, and adaptive machine learning-augmented rule engines. Implementation patterns are examined for hybrid mainframe environments, including Change Data Capture integration, synchronous pre-authorization screening, and federated cross-border deployment. Applications across Anti-Money Laundering (AML), Customer Due Diligence (CDD), and FATCA reporting domains are analyzed. The ARA is positioned as an architecture-level reference model — distinct from vendor-specific RegTech products and high-level regulatory guidance — addressing a demonstrable gap in compliance modernization literature for institutions with significant legacy infrastructure investment.

Keywords: *Adaptive Rule Engine, Anti-Money Laundering, Event-Driven Architecture, Financial Transaction Monitoring, Real-Time Compliance, RegTech*

1. Introduction

Since the passage of the Bank Secrecy Act of 1970, financial institutions have carried a legal obligation to monitor transactions for compliance with anti-money laundering, sanctions, and reporting requirements. What has changed — and changed decisively over the past decade — is the technical standard against which that obligation is now measured. Regulatory bodies across major jurisdictions, including the Financial Crimes Enforcement Network (FinCEN), the Securities and Exchange Commission (SEC), and the European Banking Authority (EBA), have progressively tightened expectations around detection latency, reporting timeliness, and the depth of entity-level due diligence that institutions must maintain [5][6][7]. Continuous transaction monitoring, near-real-time suspicious activity identification, and immediate corrective action capability are no longer aspirational standards — they are explicit regulatory requirements backed by enforcement consequences. Against this backdrop, the dominant compliance infrastructure at most established financial institutions reflects a different era's assumptions. Transaction records accumulate in relational database staging areas; compliance rule engines execute scheduled batch jobs against those records during off-peak windows; findings are reported through periodic submission cycles. This architecture was not poorly designed for its time — it was well-matched to a regulatory environment characterized by periodic inspection and financial crime techniques that operated on timescales compatible with daily batch detection. Neither of those conditions continues to hold [1][2].

The financial crime methodologies that batch architectures cannot reliably intercept are precisely

[1]*Acharya Nagarjuna University, India*
vanaparthin.rao@gmail.com
[2]*University of Central Missouri, USA*
mahesh.dhanekula9@gmail.com
[3]*VTU, Belgaum, India*
srishabh.srivastava@gmail.com

those that dominate current enforcement caseloads: structured transactions engineered to remain below individual reporting thresholds, rapid fund layering through multi-hop correspondent banking chains, and beneficial ownership concealment through shell entity networks that defeat name-matching heuristics [1][9]. Each technique exploits the detection latency window that batch compliance systems structurally create — the interval between transaction execution and compliance evaluation during which funds move and forensic context degrades.

The institutional response has characteristically involved adding RegTech capabilities atop existing infrastructure rather than replacing it. Wholesale replacement of core transaction processing systems at Tier-1 institutions carries operational and regulatory risk that most boards and regulators regard as disproportionate to the compliance improvement it might deliver [4]. The result, predictably, is a fragmented architecture: streaming compliance components imperfectly synchronized with legacy transaction records, inconsistent entity schemas generating avoidable false positives in sanctions screening, and compliance evaluation gaps at the integration boundaries between old and new systems.

This article argues that a structured architectural solution to this fragmentation is available — and that it does not require core system displacement. The Adaptive RegTech Architecture (ARA), presented here, provides a reference model for continuous, low-latency compliance evaluation across AML, CDD, and FATCA domains, designed specifically for institutions whose infrastructure spans mainframe and modern computing environments. Four design principles — compliance decoupling, event primacy, unified identity, and adaptive intelligence — constitute the framework's defining commitments [2][3].

2. Method

2.1 Structural Limitations of Legacy Compliance Architectures

Three technically distinct failure modes characterize batch-processing compliance architectures, and each compounds the others in practice. Temporal failure is the most visible: accumulating transactions before evaluation creates a detection window spanning twelve to twenty-four hours within which suspicious activity continues, transferred funds move across additional correspondent hops, and the contextual evidence needed for Suspicious Activity Report (SAR) preparation is progressively degraded [1]. For an institution processing one million daily transactions through a nightly batch cycle, the compliance evaluation system examines each transaction not against its real-time behavioral context but against a static record from which temporal ordering and inter-transaction relationships have been removed.

Contextual failure is less often discussed but equally damaging. Sanctions screening that operates against decontextualized flat records — matching transaction party names against list entries without traversing ownership structures — generates both false positives that impose operational burden and false negatives that create regulatory exposure [8][12]. Effective beneficial ownership resolution requires navigating multi-level corporate structures in real time; relational database schemas designed for transaction processing, not relationship traversal, cannot support this requirement without significant query latency penalties that render real-time application impractical.

Architectural coupling constitutes the third failure mode. Compliance rule engines that read directly from transaction databases inherit the operational constraints of those databases — they cannot scale independently of transaction volume, cannot be updated without coordination across system boundaries, and cannot apply consistent evaluation logic across source systems with heterogeneous data schemas [2][11]. Each of these constraints becomes more damaging as compliance obligations grow in scope and specificity. Table 1 maps each failure mode to its technical root cause and regulatory risk consequence.

Table 1. Structural Limitations of Legacy Compliance Architectures

Failure Mode	Technical Root Cause	Regulatory Risk Consequence
Temporal (Detection Latency)	Batch accumulation creates 12–24 hour window before compliance evaluation	Suspicious activity proceeds undetected; funds traverse additional hops before flagging
Contextual (Entity Isolation)	Flat relational records strip inter-transaction and ownership relationships	False positives from name-matching; false negatives from unresolved beneficial ownership
Architectural Coupling	Rule engines read directly from transaction databases, inheriting their operational constraints	Cannot scale independently; updates require cross-system coordination; inconsistent logic across source systems

2.2 The Adaptive RegTech Architecture — Design Principles

Each identified failure mode maps directly to one of the ARA's four design principles, establishing a logical architectural derivation rather than an asserted framework.

Temporal failure is addressed by Compliance Decoupling: separating the compliance evaluation layer from the transaction execution layer through event-driven interfaces eliminates the batch accumulation step that creates detection latency. Compliance components that consume from an event stream evaluate each transaction as it occurs, not hours after the fact [11]. Decoupling also enables compliance infrastructure to scale independently of transaction processing systems — a necessary property as transaction volumes grow faster than compliance staffing.

Contextual failure is addressed by Unified Identity and Event Primacy operating together. Event records that carry full transactional context at the moment of occurrence — party identifiers, geographic metadata, network routing information — preserve the situational detail that batch accumulation destroys [11]. A continuously maintained graph-structured entity model enables beneficial ownership resolution, relationship traversal, and network pattern detection at query time, transforming sanctions screening from a name-matching exercise into a relationship analysis capability [12][14].

Architectural coupling is addressed by Adaptive Intelligence: rule engines that operate against the event stream through the Compliance Evaluation Engine's (CEE) evaluation tier can be updated, versioned, and scaled without modifying the source transaction systems that emit events [13][15]. Machine learning models trained on annotated historical compliance outcomes — SAR filings, enforcement actions, confirmed fraud cases — detect statistical patterns that resist explicit rule formulation, extending detection coverage to emerging typologies as they appear in operational data [13][16].

2.3 ARA Component Architecture

Three processing tiers implement the ARA's design principles in sequence. The Event Streaming Layer captures transaction events from all source systems — payments, securities trades, foreign exchange settlements, loan disbursements — as immutable, schema-versioned records published to a centralized event bus. Microsecond-precision timestamps and full party and routing metadata are embedded in each record at capture. Exactly-once delivery semantics prevent compliance gaps from duplicate suppression or event loss, and retention configurations align with the five-to-seven-year record-keeping requirements common across major regulatory jurisdictions [7].

The Entity Resolution Layer maintains the graph-structured entity model that transforms individual transaction evaluation into relationship-aware compliance assessment. Graph databases provide constant-time relationship traversal regardless of graph depth — a property critical for real-time beneficial ownership resolution, where the number of ownership layers between a transacting party and

its ultimate controller varies unpredictably [12][14]. The graph is continuously refreshed from onboarding records, regulatory filings, commercial beneficial ownership data, and transaction history, ensuring compliance evaluations reflect current entity state rather than stale customer master records.

The Compliance Evaluation Engine operates in two tiers that execute in parallel against each event. Deterministic rules encode regulatory thresholds, sanctions list matching with entity graph traversal for beneficial ownership resolution, and product-specific compliance logic — producing auditable, explainable outputs for each evaluation. The ML tier operates behavioral baseline models trained on entity-specific historical patterns across transaction frequency, amount distributions, geographic counterparty dispersion, and product utilization [13][15]. The CEE aggregates outputs from both tiers into composite risk signals routed to clear, review, escalation, or block dispositions.

3. Results and Discussion

3.1 Integration Patterns for Hybrid Mainframe Environments

Change Data Capture (CDC) is the integration mechanism that makes the ARA compatible with mainframe-resident transaction processing cores without requiring those cores to be modified or replaced. CDC monitors database transaction logs and emits structured change events for every committed transaction, publishing them to the ARA's event streaming layer with sub-second latency from commit to publication [4]. This non-invasive pattern has been applied in production across the range of mainframe database environments that characterize established financial institutions — IBM Db2 on z/OS, IMS DB, Oracle mid-tier configurations — environments whose operational characteristics are directly familiar from the author's delivery experience across Tier-1 financial institutions, including AML, CDD, and FATCA compliance system delivery at HSBC and securities lending transaction processing at State Street Corporation.

The primary implementation discipline is schema governance: CDC configurations must capture the full semantic payload required for compliance evaluation, not only the fields required for transaction processing. As core system data models evolve through application changes, CDC configurations must be maintained in alignment — a governance requirement that, when neglected, produces silent compliance data quality degradation without triggering operational alerts. CDC implementations targeting compliance data pipelines achieve sub-second latency from transaction commit to event publication, satisfying the ARA's real-time processing requirements across all compliance domains.

Compliance domains requiring pre-authorization — real-time sanctions screening before payment execution — demand a synchronous evaluation pattern in which the transaction processing system queries the CEE through a low-latency API before committing the transaction. Managing this pattern within consumer-facing response time budgets requires careful optimization of entity graph traversal performance and CEE rule evaluation throughput. Production deployments of this pattern have achieved end-to-end compliance evaluation latencies of 40 to 150 milliseconds for standard transactions, with elevated latency for high-complexity beneficial ownership structures that require deeper graph traversal [11][12]. Table 2 summarizes integration patterns and latency profiles across the principal hybrid deployment scenarios.

Table 2. ARA Component Integration Patterns and Latency Profiles

Integration Pattern	Source Environment	Latency Profile	Compliance Use Case
Change Data Capture (Async)	IBM Db2 z/OS, IMS DB, Oracle	Sub-second (commit to event publication)	AML monitoring, CDD re-scoring, FATCA identification
Synchronous Pre-Authorization	All mainframe and mid-tier cores via CEE API	40–150 ms (standard transactions)	Real-time sanctions screening before payment execution

Federated Regional Deployment	Multi-jurisdiction cloud-hosted regional instances	Regional: <100 ms; cross-border signal aggregation: <500 ms	Cross-border AML and data residency-constrained reporting
Event Replay (Historical Analysis)	Immutable event log with 5–7 year retention	Batch (minutes to hours depending on scope)	Regulatory examination support, SAR preparation, litigation discovery

3.2 Regulatory Domain Applications

The AML detection architecture within the ARA departs from the alert-per-transaction model of legacy monitoring systems. Each customer entity maintains continuously updated behavioral baseline models across multiple transaction dimensions — frequency, amount variance, counterparty geographic spread, and product utilization patterns. Deviations from these baselines, weighted by entity risk tier, generate composite anomaly scores rather than binary alert flags, enabling prioritization of investigative resources toward the highest-risk signals rather than uniform treatment of all threshold crossings [13][16].

Graph-based detection extends AML coverage to network-level patterns that individual transaction monitoring cannot identify. Ring-structuring schemes — coordinated deposits across multiple entities, each below individual reporting thresholds, collectively constituting a reportable aggregate — are detectable through graph analysis of transaction relationships across the entity network [14][17]. Multi-hop layering through correspondent banking chains requires traversing transaction relationships across multiple graph nodes; the ARA's entity resolution layer makes this traversal a standard query operation rather than a multi-system data assembly exercise requiring manual investigation hours [17][18].

CDD obligations present a different architectural challenge: maintaining current customer risk profiles against a continuously changing external information environment. The ARA addresses this through continuous integration of external data sources — commercial registry updates, adverse media feeds, politically exposed person lists, and beneficial ownership registry changes — into the entity graph, with automated re-assessment workflows triggered by material changes [8][18]. Customers crossing defined risk thresholds are queued for enhanced due diligence review with entity graph state and transaction history pre-assembled, addressing the analyst throughput constraints that limit conventional CDD program effectiveness.

FATCA identification illustrates the ARA's capacity to convert a complex, historically manual reporting obligation into a continuously maintained automated process. Beneficial ownership chain traversal, applied continuously as new account and entity data arrives, identifies accounts requiring FATCA classification without periodic batch re-screening [10]. US indicia detection — addresses, telephone numbers, birthplace records, and fund transfer destination patterns — operates against each new customer interaction, capturing classification-relevant information at the point of origin rather than through retrospective data review [10][7]. IRS Form 8966 data elements are pre-populated from entity graph state, materially reducing the manual assembly work that creates reporting cycle bottlenecks in institutions processing FATCA obligations through conventional methods.

3.3 Performance, Scalability, and Cross-Border Considerations

Horizontal scalability is a structural property of the ARA's event streaming foundation, not a capability requiring architectural modification as transaction volumes grow. Consumer instances in the compliance evaluation pipeline are stateless with respect to the event stream — each instance reads from a partition of the event log and processes its assigned events independently. Adding instances to handle increased transaction volume requires no changes to compliance logic or the entity graph; near-linear throughput scaling characterizes well-implemented streaming architectures, enabling proportional capacity expansion without architectural redesign [11].

Entity resolution performance under high transaction volumes requires active management. Deep beneficial ownership graph traversal has variable computational cost depending on

ownership structure complexity. Pre-computed ownership chain summaries for frequently accessed entities and caching strategies for high-frequency counterparty lookups — major correspondent banks and institutional trading counterparties — can reduce effective graph query latency by sixty to eighty percent in typical institutional portfolios, converting deep traversal operations into cache-hit retrievals for the majority of transaction volume [12].

Multi-jurisdiction deployments introduce data residency constraints that cannot be resolved through technical optimization alone. The federated ARA pattern deploys regional compliance evaluation instances within each jurisdiction, each operating against a regional entity graph containing only locally permitted data. Cross-border transaction analysis proceeds through anonymized risk signal aggregation rather than raw data sharing, preserving compliance analytical capability without creating data sovereignty violations across the regulatory environments that govern Tier-1 international financial institutions [19][4].

3.4 Societal and Institutional Implications

The case for real-time compliance infrastructure is routinely made in terms of institutional risk reduction — regulatory penalty avoidance, reputational protection, and operational efficiency gains. It is less frequently framed as a contribution to financial system integrity and social equity, though the evidence for both dimensions is substantial. Reducing detection latency windows removes the operational space within which financial crime networks function; money laundering operations that finance trafficking, narcotics distribution, and terrorism financing depend on the ability to move funds through the financial system faster than compliance monitoring can identify and flag the transactions. The United Nations Office on Drugs and Crime estimates that approximately two to five percent of global GDP is laundered annually — a figure that improved compliance detection infrastructure has measurable potential to reduce over sustained deployment timeframes [9].

Algorithmic compliance systems introduce a distinct category of institutional obligation. ML-based risk scoring trained on historical compliance data may encode the demographic and geographic biases present in that data, producing false-positive rates that fall disproportionately on specific customer segments [19][15]. Institutions that deploy ARA-style architectures without monitoring deployed models for disparate impact are trading one regulatory risk for another — the risk of discriminatory outcomes in compliance decisions that affect account access, transaction processing, and SAR filing rates. Explainability interfaces that allow compliance officers to interrogate model-generated risk scores are both an internal governance requirement and an increasingly codified regulatory expectation in jurisdictions examining algorithmic compliance systems [15][20].

Financial de-risking — the practice of exiting customer segments, geographies, or product categories perceived as carrying elevated compliance cost — is a well-documented consequence of high false-positive rates in compliance monitoring [18]. Immigrant communities relying on remittance services, small businesses in cash-intensive sectors, and customers in jurisdictions with elevated risk ratings bear the primary cost of de-risking decisions driven by imprecise compliance detection. By reducing false-positive rates while sustaining detection sensitivity, the ARA addresses a false economy that purely reactive compliance architectures perpetuate: the compliance cost that institutions appear to eliminate by exiting high-risk segments is real, but so is the regulatory exposure created by the detection blind spots that segment exit cannot eliminate.

4. Conclusion

The regulatory compliance obligations that financial institutions carry have outgrown the batch-processing architectures through which most institutions discharge them. Detection latency windows of twelve to twenty-four hours, compliance evaluation divorced from real-time transaction context, and rule engines architecturally coupled to the systems they monitor are not incidental limitations of aging technology — they are structural properties of a compliance paradigm that was not designed for continuous, relationship-aware, adaptive evaluation against the regulatory mandates and financial crime methodologies that now define the operating environment.

What the Adaptive RegTech Architecture offers is not a technology replacement but an architectural reorientation organized around four design principles that reposition compliance evaluation as

an independent, continuously operating capability. Compliance decoupling eliminates the batch accumulation step and enables independent scaling. Event primacy preserves the transactional context that batch processing destroys. Unified identity makes beneficial ownership resolution and network relationship analysis real-time capabilities rather than multi-day investigation exercises. Adaptive intelligence extends detection coverage to evolving typologies that static rule sets cannot anticipate.

These principles are achievable within the hybrid infrastructure environments that define most established financial institutions. CDC integration connects mainframe transaction cores to the ARA's event streaming layer without modification or displacement of those cores. Pre-authorization screening patterns operate within the latency budgets that mainframe-adjacent transaction processing requires. Federated deployment patterns accommodate the data residency constraints that multi-jurisdiction operations impose. The path from fragmented legacy compliance to architecturally coherent real-time compliance does not require institutions to first solve the problem of core system replacement — it requires drawing the compliance layer boundary in the right place and connecting it to existing infrastructure through interfaces that preserve the auditability, transactional integrity, and regulatory validation history that mainframe systems carry. Institutions that implement the ARA systematically are building detection infrastructure capable of operating at the pace of modern financial crime and contributing, in a measurable way, to the integrity of the financial system on which their customers and counterparties depend.

References

[1] D. W. Arner, J. Barberis, and R. P. Buckley, "FinTech, RegTech, and the reconceptualization of financial regulation," Northwestern Journal of International Law and Business, vol. 37, no. 3, pp. 371–413, 2016. https://scholarlycommons.law.northwestern.edu/njilb/vol37/iss3/2/

[2] D. A. Zetzsche, R. P. Buckley, D. W. Arner, and J. N. Barberis, "Regulating a revolution: From regulatory sandboxes to smart regulation," Fordham Journal of Corporate and Financial Law, vol. 23, no. 1, pp. 31–103, 2018. https://ir.lawnet.fordham.edu/jcfl/vol23/iss1/2/

[3] C. M. Kahn and W. Roberds, "Why pay? An introduction to payments economics," Journal of Financial Intermediation, vol. 18, no. 1, pp. 1–23, 2009. https://www.sciencedirect.com/science/article/abs/pii/S1042957308000533

[4] Basel Committee on Banking Supervision, "Sound practices for the management and supervision of operational risk," Bank for International Settlements, Basel, Switzerland, 2021. https://www.bis.org/publ/bcbs183.pdf

[5] Financial Crimes Enforcement Network, "Anti-money laundering program and suspicious activity report filing requirements for registered investment advisers," U.S. Department of the Treasury, Washington, DC, 2015. https://www.federalregister.gov/documents/2015/09/01/2015-21318/anti-money-laundering-program-and-suspicious-activity-report-filing-requirements-for-registered

[6] European Banking Authority, "Guidelines on internal governance under Directive 2013/36/EU," EBA/GL/2021/05, European Banking Authority, Paris, France, 2017. https://www.bankofengland.co.uk/-/media/boe/files/paper/2020/december/guidelines-on-internal-governance---revised.pdf

[7] "Electronic Recordkeeping Requirements for Broker-Dealers, Security-Based Swap Dealers, and Major Security-Based Swap Participants," Federal Register, Nov. 03, 2022. https://www.federalregister.gov/documents/2022/11/03/2022-22670/electronic-recordkeeping-requirements-for-broker-dealers-security-based-swap-dealers-and-major

[8] E. Curry, "Trusted digital identities in financial services: new MLR guidance announced – Enabling digital identity," Blog.gov.uk, Jul. 23, 2025. https://enablingdigitalidentity.blog.gov.uk/2025/07/23/trusted-digital-identities-in-financial-services-new-mlr-guidance-announced/

[9] UNODC, "World Drug Report 2022," United Nations : Office on Drugs and Crime, 2022. https://www.unodc.org/unodc/en/data-and-analysis/world-drug-report-2022.html

[10] "FATCA information for foreign financial institutions and entities," Internal Revenue Service, https://www.irs.gov/businesses/corporations/information-for-foreign-financial-institutions

[11] P. Kumar Mantha, "Event-Driven Architectures in Financial Services." IJIRMPS. [Online]. Available: https://www.ijirmps.org/papers/2019/3/232941.pdf.

[12] E. Kurshan and H. Shen, “Graph Computing for Financial Crime and Fraud Detection: Trends, Challenges and Outlook,” International Journal of Semantic Computing, vol. 14, no. 04, pp. 565–589, Dec. 2020, doi: https://doi.org/10.1142/s1793351x20300022.

[13] W. Bi, T. K. Trinh, and S. Fan, “Machine Learning-Based Pattern Recognition for Anti-Money Laundering in Banking Systems,” Journal of Advanced Computing Systems, vol. 4, no. 11, pp. 30–41, Nov. 2024, doi: https://doi.org/10.69987/jacs.2024.41103.

[14] I. Alarab, S. Prakoonwit, and M. I. Nacer, “Competence of Graph Convolutional Networks for Anti-Money Laundering in Bitcoin Blockchain,” Proceedings of the 2020 5th International Conference on Machine Learning Technologies, Jun. 2020, doi: https://doi.org/10.1145/3409073.3409080.

[15] D. Cheng, Y. Zou, S. Xiang, and C. Jiang, “Graph Neural Networks for Financial Fraud Detection: A Review,” arXiv (Cornell University), Oct. 2024, doi: https://doi.org/10.1007/s11704-024-40474-y.

[16] V. Van Vlasselaer et al., “APATE: A novel approach for automated credit card transaction fraud detection using network-based extensions,” Decision Support Systems, vol. 75, pp. 38–48, Jul. 2015, doi: https://doi.org/10.1016/j.dss.2015.04.013.

[17] F. Colladon and E. Remondi, "Using social network analysis to prevent money laundering," Expert Systems with Applications, vol. 67, pp. 49–58, 2021. DOI: https://ui.adsabs.harvard.edu/abs/2021arXiv210505793F/abstract

[18] E. Tsingou, “Global financial governance and the developing anti-money laundering regime: What lessons for International Political Economy?,” International Politics, vol. 47, no. 6, pp. 617–637, Nov. 2010, doi: https://doi.org/10.1057/ip.2010.32.

[19] H. Zhang, J. Hong, D. Fan, S. Drew, L. Xue, and J. Zhou, “A Privacy-Preserving Hybrid Federated Learning Framework for Financial Crime Detection,” arXiv (Cornell University), Apr. 2023, doi: https://doi.org/10.48550/arxiv.2302.03654.

[20] S. K. Parimi and R. Yallavula, “Data-Governed Autonomous Decisioning: AI Models for Real-Time Optimization of Enterprise Financial Journeys,” International Journal of Emerging Trends in Computer Science and Information Technology, vol. 2, 2021, doi: https://doi.org/10.63282/3050-9246.ijetcsit-v2i1p111.