Sarcouncil Journal of Multidisciplinary

ISSN(Online): 2945-3445

Volume- 05| Issue- 07| 2025

Review Article **Received:** 06-06-2025 | **Accepted:** 29-06-2025 | **Published:** 19-07-2025

Real-Time Identity Resolution at Scale: Technical Implementation Guide

Tejendra Patel
California State University, Los Angeles (CSULA), CA, USA

Abstract: Real-time identification resolution at agency scale represents a crucial technological capability that addresses the growing complexity of virtual user identity throughout multiple systems and gadgets. Cutting-edge organizations face unprecedented challenges in preserving correct consumer profiles as individuals engage via various touchpoints, create multiple accounts, and make use of various payment strategies across digital ecosystems. Graph-based, totally factual architectures offer the foundational framework for effective identification correlation via representing entities and relationships as interconnected nodes, allowing sophisticated analysis of consumer connections across gadgets, bills, and behavioral styles. Superior streaming structures facilitate high-performance statistics ingestion pipelines able to process hundreds of thousands of activities consistently with 2nd whilst retaining sub-2nd latency necessities crucial for real-time decision making. Probabilistic matching algorithms generate graduated self-perception scores that mirror healthy truth based on signal energy and historical accuracy, moving beyond binary matching selections to provide nuanced entity selection abilities. Statistics consistency mechanisms put in force eventual consistency patterns with state-of-the-art war resolution strategies that prioritize record freshness and temporal relevance throughout distributed processing nodes. The combination of identity resolution talents with fraud detection systems enables contextual enforcement actions along with chance-based total authentication and device management strategies. Employer implementations exhibit sizeable upgrades in fraud prevention effectiveness whilst preserving the most reliable person reviews via sensible authentication mechanisms that adapt dynamically to behavioral styles and threat tests.

Keywords: Identity resolution, graph-based architecture, probabilistic matching, real-time processing, fraud detection, adaptive authentication.

INTRODUCTION

Cutting-edge digital platforms face an increasing number of complex undertakings: appropriately figuring out and tracking customers throughout multiple touchpoints in real-time. As customers engage through numerous devices, create multiple accounts, and make use of exclusive price techniques, conventional identification control methods fall quickly. Modern virtual ecosystems witness exceptional complexity in user authentication and verification processes, in which agencies struggle to hold coherent person profiles throughout fragmented interaction channels. The proliferation of virtual touchpoints has basically transformed how companies approach identification control, growing scenarios in which single individuals preserve disparate virtual personas across a couple of platforms and services (Viswanathan, L. 2024).

Virtual identification control systems are subject to sizable threats from state-of-the-art attack vectors, including credential stuffing, account takeover attempts, and synthetic identity fraud. Business enterprise environments revel in increased vulnerability levels as cybercriminals take advantage of weaknesses in conventional authentication mechanisms, especially concentrated on corporations with legacy identity infrastructure. The panorama of digital identification threats keeps evolving rapidly, with rising assault styles leveraging synthetic intelligence and machine learning strategies to bypass traditional security features. Corporations face mounting stress in enforcing comprehensive identification solutions that can adapt to dynamic threat environments, even as they retain operational efficiency and consumer enjoyment standards (Viswanathan, L. 2024).

Real-time identity resolution has emerged as a vital functionality for fraud prevention, personalization, and operational performance. Patron identification decision structures require sophisticated architectural frameworks capable of processing big data volumes while maintaining accuracy in entity matching and courting dedication. The complexity of present-day identity resolution extends past simple consumer identity, encompassing complete purchaser journey mapping throughout multiple interaction channels and touchpoints. Enterprise-scale identification resolution structures ought to manage numerous data sources consisting of transactional data, behavioral patterns, tool fingerprints, and demographic records to assemble accurate patron profiles (Garzon, C. 2024).

Technical exploration examines the engineering standards and architectural styles required to construct strong identity decision structures that are able to process excessive-speed information streams, even while preserving accuracy and consistency at an agency scale. Implementation of effective customer identification decisions needs cautious attention to records ingestion pipelines,

***Corresponding Author:** Tejendra Patel

DOI- https://doi.org/10.5281/zenodo.16149020

entity matching algorithms, and conflict resolution mechanisms. Modern-day identification systems have to accommodate both batch processing necessities for ancient information evaluation and real-time processing capabilities for immediate fraud detection and personalization efforts. The architectural complexity will increase when thinking about the need for scalable graph-based record models, which could efficiently constitute and question complicated customer dating networks (Garzon, C. 2024).

The evolution toward complete identity resolution reflects broader industry popularity that traditional processes can't effectively cope with the current virtual demanding situations. Organizations imposing advanced identity control measures need to balance safety necessities with user experience concerns while complying with evolving privacy regulations. Success in cutting-edge identification decision requires integration of more than one technological additive consisting of device learning algorithms, allocated processing frameworks, and advanced analytics competencies that could perform cohesively within company generation stacks.

GRAPH-BASED DATA ARCHITECTURE FOR IDENTITY CORRELATION

The foundation of effective identification decisions lies in adopting graph-based statistical models that represent entities and relationships as interconnected nodes. In contrast to conventional relational databases, graph systems seize the complex internet of connections between users, gadgets, bills, and behavioral alerts. Identity graphs essentially transform how organizations manage patron records by developing unified perspectives of individual clients across multiple touchpoints and interaction channels. Contemporary identity graph implementations leverage sophisticated data systems that may efficiently represent millions of client relationships while retaining query overall performance appropriate for real-time applications (Patino, A. 2025).

Modern-day identification graph architectures exhibit first-rate scalability traits, with leading implementations capable of processing billions of patron interactions while maintaining sub-2nd question response times. The architectural benefit of graph-based procedures will become especially glaring while coping with complex customer experience evaluation that requires traversing multiple data types concurrently. Company-scale identity graphs typically encompass various records assets, including transactional statistics, behavioral analytics, demographic facts, and engagement metrics to assemble complete customer profiles. Performance optimization techniques permit identification of graphs to help concurrent queries from a couple of business applications without degrading device responsiveness (Patino, A. 2025).

Core Entity Types and Relationships

The graph model encompasses several primary entity sorts: person profiles, device fingerprints, conversation endpoints such as e-mail addresses and phone numbers, payment instruments, and behavioral styles. Every entity maintains attributes together with advent timestamps, verification repute, and confidence scores that permit state-of-the-art matching algorithms. Multi-modal search abilities within knowledge graphs enable sophisticated entity resolution through combining textual, numerical, and categorical statistical sources to identify client relationships with excessive accuracy. Superior graph architectures guide complicated query styles that can simultaneously examine more than one entity attribute, even as they retain computational performance across massive-scale datasets (Hypermode, 2025).

Tool fingerprinting bureaucracy is an important issue, combining browser traits, display screen decision, timezone statistics, and hardware specs to create specific tool signatures. Present-day fingerprinting tactics contain a couple of statistical modalities consisting of behavioral patterns, hardware configurations, and software program environments to establish chronic device identity. Know-how graph implementations leverage multi-modal seek techniques to correlate tool fingerprints across distinct statistics resources, enabling accurate tool identification even if partial fingerprint records are available. Superior fingerprinting methodologies gain excellent staying power fees, with tool signatures retaining balance across extended time durations and more than one consumer session (Hypermode, 2025).

Communique endpoints serve as sturdy identification anchors, especially when demonstrated through multi-component authentication tactics. Email addresses and phone numbers offer dependable correlation vectors that enable pass-platform client identification with high confidence levels. Multi-modal search methods within identification graphs can successfully

Publisher: SARC Publisher

integrate communication endpoint facts with behavioral styles and transaction histories to set up complete customer profiles. The mixing of tested communication endpoints with extra identification alerts creates robust client identity mechanisms that aid both fraud prevention and personalization programs (Patino, A. 2025).

Payment gadgets offer another reliable correlation vector, as customers normally hold steady payment strategies throughout debts and platforms. Monetary records integration within identity graphs requires cautious attention to privacy rules and protection requirements, even as permitting effective purchaser identity. Know-how graphs excel at correlating fee styles with client behavior across a couple of dimensions, helping state-of-the-art fraud detection and hazard evaluation packages. The aggregate of price tool records with other identification signals creates effective correlation mechanisms that permit correct customer identity while retaining compliance with economic privacy regulations (Hypermode, 2025).

REAL-TIME DATA PROCESSING AND INGESTION

High-overall performance identity resolution demands state-of-the-art statistics ingestion pipelines capable of processing tens of millions of events per second with sub-second latency requirements. The structure employs allotted streaming structures to deal with real-time person interactions, login events, transaction information, and device signals. Apache Kafka represents the enterprise trend for constructing strong information pipelines that can take care of big throughput requirements while maintaining record integrity and fault tolerance. Present-day Kafka implementations display incredible scalability characteristics, with production clusters assisting throughput stages exceeding several million messages per second across thousands of topics and partitions. The architectural basis of high-throughput streaming structures requires cautious attention to community topology, garage infrastructure, and computational resources to gain top-quality performance under varying load situations (Ethan, 2024).

Agency-scale streaming structures need to accommodate numerous information resources such as application logs, database exchange streams, IoT sensor records, and user interaction activities. Kafka's distributed architecture allows horizontal scaling throughout more than one dealer node, with every broker capable of managing heaps of concurrent connections while preserving message ordering guarantees within a single cluster. Superior partitioning strategies ensure that associated occasions are always routed to the same processing nodes, permitting stateful processing operations that require retaining context across a couple of messages. The combination of producer batching, compression algorithms, and optimized serialization codecs contributes to accomplishing most throughput whilst minimizing community bandwidth intake (Ethan, 2024).

Stream Processing Architecture

Event-pushed processing pipelines consume statistics from more than one source simultaneously, making use of initial filtering and enrichment before feeding into the identification resolution engine. Apache Kafka serves as the backbone for dependable message shipping, with custom partitioning strategies making sure related events attain the same processing nodes for constant data management. Provider stage settlement (sla) primarily based model schemes play a vital role in maintaining consistent overall performance throughout distributed movement processing engines. Studies demonstrate that adaptive resource allocation mechanisms can enhance device throughput by up to 35% while decreasing latency violations by about forty in comparison to static aid allocation methods (Hanif, M. *et al*., 2019).

The system implements windowing techniques to batch associated events occurring within particular timeframes, lowering computational overhead and preserving near real-time responsiveness. Sla-aware model algorithms continuously screen machine performance metrics and automatically alter resource allocation to meet predefined performance objectives. Experimental outcomes suggest that dynamic version schemes can preserve sla compliance charges above ninety % even underneath extraordinarily variable workload conditions. Advanced model mechanisms contain predictive modeling strategies that assume aid necessities based on historical overall performance statistics and contemporary system nation (Hanif, M. *et al*., 2019).

Dynamic scaling mechanisms mechanically alter processing capability based on incoming data pace, preventing bottlenecks during traffic spikes. Kafka's producer and customer configurations may be optimized to achieve the highest throughput, with the right tuning allowing individual producers to acquire fees exceeding 100,000 messages consistent with 2nd. Buffer control strategies and

Publisher: SARC Publisher

asynchronous processing patterns make substantial contributions to overall system performance, with properly configured structures demonstrating latency characteristics appropriate for real-time programs. The implementation of effective tracking and alerting systems guarantees that overall performance degradation may be detected and addressed before sla violations occur (Ethan, 2024).

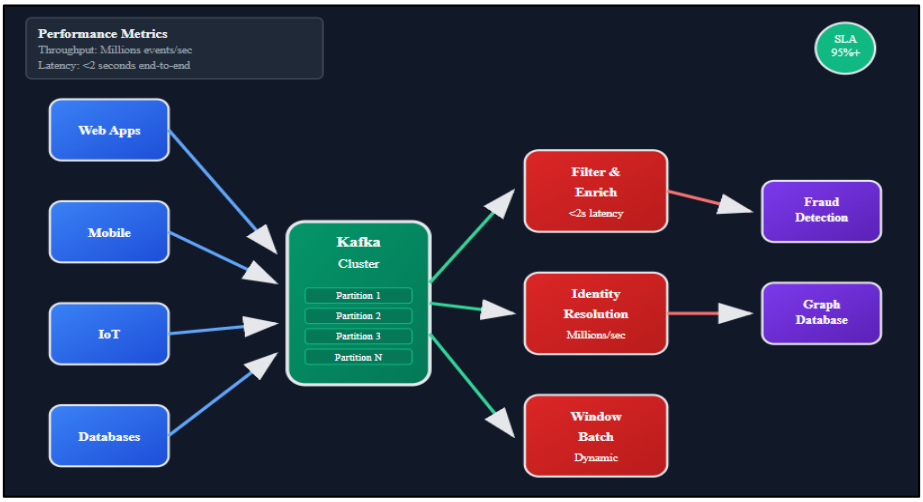

Fig 1. Real-Time Data Processing Pipeline (Ethan, 2024; Hanif, M. *et al*., 2019).

ENTITY RESOLUTION LOGIC AND CONFIDENCE SCORING

The middle decision engine applies probabilistic matching algorithms to determine entity relationships with quantified self-perception tiers. In the choice of binary matching choices, the gadget generates graduated self-belief rankings reflecting the shape of reality based on sign power and historical accuracy. Probabilistic matching represents a massive advancement over conventional fuzzy matching strategies by incorporating statistical models that can cope with uncertainty and provide more satisfactory versions more effectively. Not like fuzzy matching, which relies more often than not on string similarity algorithms, probabilistic matching evaluates more than one statistical attribute simultaneously while considering the probability of matches based on comprehensive statistical analysis. Modern-day probabilistic matching structures reveal superior accuracy in managing real-world fact-demanding situations along with typographical mistakes, missing data, and formatting inconsistencies (Novoselsky, B. 2024).

Multi-Signal Matching Framework

The matching procedure evaluates a couple of indicators simultaneously, applying weighted scoring primarily based on signal reliability and expertise. Email addresses and validated phone numbers receive better weights than behavioral styles or device traits. Temporal elements affect scoring, with the latest shared signals carrying more weight than historical connections. Entity resolution encompasses twelve number one strategies, which include actual matching, fuzzy matching, phonetic matching, and machine learning-based strategies that can be mixed to achieve the best outcomes across diverse datasets. Studies demonstrate that combining multiple matching strategies can appreciably enhance standard accuracy in comparison to single-approach processes, with ensemble techniques showing particular effectiveness in handling complicated entity decision challenges (Otten, N. V. 2024).

Device learning models continuously refine self-assurance thresholds primarily based on comment loops from guide reviews and downstream validation methods. The system maintains separate confidence fashions for special use instances, spotting that fraud detection calls for higher precision than advertising personalization efforts. Superior entity decision implementations leverage Python and R libraries that provide state-of-the-art matching abilities such as probabilistic report linkage, blocking algorithms, and similarity computation features. The provision of specialised libraries allows agencies to put in force complicated entity decision workflows without developing custom algorithms from scratch, significantly decreasing development time and improving reliability (Otten, N. V. 2024).

Handling Ambiguous Cases

Facet cases consisting of shared gadgets in families or public computer systems require state-

Publisher: SARC Publisher

of-the-art disambiguation logic. The gadget analyzes behavioral patterns, geographic consistency, and temporal sequences to distinguish between legitimate shared access and capability account compromise. Superior clustering algorithms are used for organizing related entities, even while maintaining person identification barriers. Entity resolution systems have to address numerous ambiguous situations, including call versions, cope with modifications, and shared contact information that can create fake advantageous fits. Blocking off techniques play an important role in handling computational complexity by reducing the range of capability suites that require certain assessment, with powerful blockading strategies attaining large overall performance enhancements even as maintaining matching accuracy (Otten, N. V. 2024).

Probabilistic matching excels in handling ambiguous cases by presenting graduated confidence scores that allow more sophisticated decision-making methods than binary matching results. The statistical foundation of probabilistic techniques permits structures to quantify uncertainty and make informed choices about borderline cases that conventional fuzzy matching can't manage effectively. Superior disambiguation strategies incorporate contextual records and area-specific information to clear up conflicts and hold data to first-class requirements throughout big-scale entity resolution operations (Novoselsky, B. 2024).

Fig 2. Entity Resolution Confidence Scoring Distribution (Novoselsky, B. 2024; Otten, N. V. 2024).

Data Consistency and Conflict Resolution

Maintaining data consistency across dispensed processing nodes affords tremendous technical challenges, mainly while dealing with conflicting information from multiple resources. The device implements eventual consistency patterns with war decision mechanisms that prioritize high-quality and recent data. Facts Disbursed systems have to navigate complicated trade-offs between consistency, availability, and partition tolerance in step with the CAP theorem, which essentially constrains gadget design choices. Agencies face crucial choices whilst choosing suitable consistency fashions, with each technique supplying extraordinary blessings and obstacles depending on specific utility necessities and operational constraints. Robust consistency models ensure on-the-spot facts coherence throughout all nodes; however, they can also sacrifice availability at some stage in community walls, whilst eventual consistency tactics prioritize machine availability at the price of transient statistics inconsistencies (Thurman, E. 2025).

The selection of suitable consistency fashions relies heavily on enterprise requirements, with financial programs normally requiring robust consistency guarantees, even as social media systems can often tolerate eventual consistency techniques. Overall performance implications vary drastically between consistency fashions, with sturdy consistency operations frequently requiring additional community round-trip journeys and coordination overhead that can impact device responsiveness. Contemporary disbursed architectures put in force hybrid techniques that combine multiple consistency fashions within the same gadget, allowing one-of-a-kind data types and operations to apply the most suitable consistency ensures (Thurman, E. 2025).

Publisher: SARC Publisher

Canonicalization and Deduplication

Standardization techniques normalize statistics codecs before storage, converting phone numbers to worldwide codecs, standardizing email addresses, and making use of constant hashing to tool fingerprints. Deduplication algorithms become aware of and merge redundant entities whilst preserving audit trails for compliance requirements. Statistics deduplication techniques encompass multiple tactics consisting of genuine matching, fuzzy matching, and gadget-studying-based strategies that could discover duplicate statistics across various formats and representations. Effective deduplication implementations usually acquire full-size garage savings, with many organizations reporting reductions of 50-90% in data garage necessities through comprehensive duplication removal approaches (Lee, S. 2025).

Real-time war decision mechanisms manage scenarios where new statistics contradict current entity relationships. The system applies versioning strategies that keep historic states while presenting regular contemporary views to downstream clients. Advanced deduplication structures comprise state-of-the-art algorithms that could take care of variations in statistics formatting, spelling mistakes, and partial matches while keeping high accuracy rates. Gadget mastering procedures for deduplication display superior overall performance as compared to traditional rule-based strategies, especially while handling complex datasets containing more than one data type and codecs. The implementation of non-stop deduplication techniques permits agencies to maintain information at high standards while processing ongoing records streams (Lee, S. 2025).

Contemporary warfare resolution frameworks need to cope with diverse technical challenges, which include network delays, node failures, and concurrent updates from more than one source. Timestamp-based ordering mechanisms provide one method to warfare resolution, even though clock synchronization across allotted nodes presents additional complexity. Vector clocks and logical timestamps provide alternative strategies that may take care of eventualities where physical clock synchronization is impractical or unreliable. The choice of conflict resolution approach significantly affects a gadget's overall performance and consistency, and requires cautious assessment of trade-offs among simplicity, accuracy, and computational overhead (Thurman, E. 2025). Deduplication accuracy relies heavily on the sophistication of similarity detection algorithms and the first-class of information preprocessing steps. Organizations enforcing comprehensive deduplication packages often study substantial improvements in statistics, great metrics, which include reduced storage expenses, advanced query performance, and more suitable analytical results via the elimination of duplicate statistics (Lee, S. 2025).

Fraud Detection and Contextual Actions

Identification resolution abilities allow sophisticated fraud detection mechanisms that leverage behavioral context and customer analysis. In place of depending solely on character transaction characteristics, the system evaluates movements within the broader context of user identification graphs and peer conduct styles. Graph analytics represents a transformative approach to fraud detection that permits companies to identify complicated fraud patterns via data analysis and network visualization. Contemporary fraud schemes often involve coordinated sports throughout multiple accounts and entities, making conventional transaction-based detection strategies insufficient for complete fraud prevention. Graph-based totally fraud detection structures excel at uncovering hidden connections among seemingly unrelated entities, revealing fraud jewelry and prepared crooked sports that would remain invisible to standard analytical tactics (NebulaGraph, 2023).

The electricity of graph analytics in fraud detection lies in the potential to research full-size networks of relationships simultaneously, allowing the identification of suspicious patterns throughout millions of interconnected information points. Monetary institutions and e-commerce platforms are increasingly relying on graph-based methods to uncover sophisticated fraud schemes such as cash laundering, artificial identification fraud, and coordinated account takeover assaults. The visible representation of the talents of graph analytics provides investigators with intuitive equipment for understanding complex fraud networks and identifying previously unknown connections among fraudulent entities (NebulaGraph, 2023).

Risk-Based Authentication

Dynamic authentication requirements alter primarily based on identification self-belief degrees and behavioral anomalies. Users with robust identity verification and steady behavioral patterns enjoy frictionless authentication, even as suspicious activities trigger extra verification steps. Tool banning mechanisms routinely block

Publisher: SARC Publisher

compromised gadgets whilst preserving valid person access via alternative devices. Adaptive authentication represents a full-scale evolution in cybersecurity generation that intelligently adjusts safety necessities based on contextual risk factors and behavioral evaluation. Unlike traditional static authentication strategies that follow uniform security measures for all customers, adaptive systems dynamically alter authentication requirements based on real-time threat evaluation and person behavior styles (Agarwal, J. 2025).

The system permits graduated reaction techniques, from passive monitoring to energetic intervention, primarily based on hazard evaluation self-belief ranges. Integration with downstream enforcement systems allows for real-time action execution without manual intervention. Current adaptive authentication implementations incorporate gadget learning algorithms that constantly examine user behavior styles to set up baseline profiles for legitimate users. When deviations from set-up behavioral norms are detected, the device routinely increases authentication necessities or implements additional security measures to verify consumer identity. Superior adaptive authentication systems can seamlessly combine a couple of authentication elements, which include biometrics, device traits, and behavioral analytics, to create complete security frameworks (Agarwal, J. 2025).

Graph-based fraud detection systems demonstrate top-notch effectiveness in identifying complicated fraud networks via superior pattern popularity and relationship analysis abilities. The ability to process and examine hundreds of thousands of transactions and relationships in real-time enables companies to locate fraud attempts before widespread economic damage takes place. Integration of graph analytics with machine learning algorithms creates powerful fraud detection engines that may adapt to evolving fraud techniques whilst retaining high accuracy rates and minimizing fake high-quality alerts (NebulaGraph, 2023).

Adaptive authentication systems offer substantial benefits over traditional safety techniques by balancing security effectiveness with consumer experience issues. Agencies enforcing adaptive authentication generally observe great upgrades in both protection results and consumer satisfaction metrics, as legitimate users experience reduced authentication friction at the same time as capability threats face elevated security scrutiny (Agarwal, J. 2025).

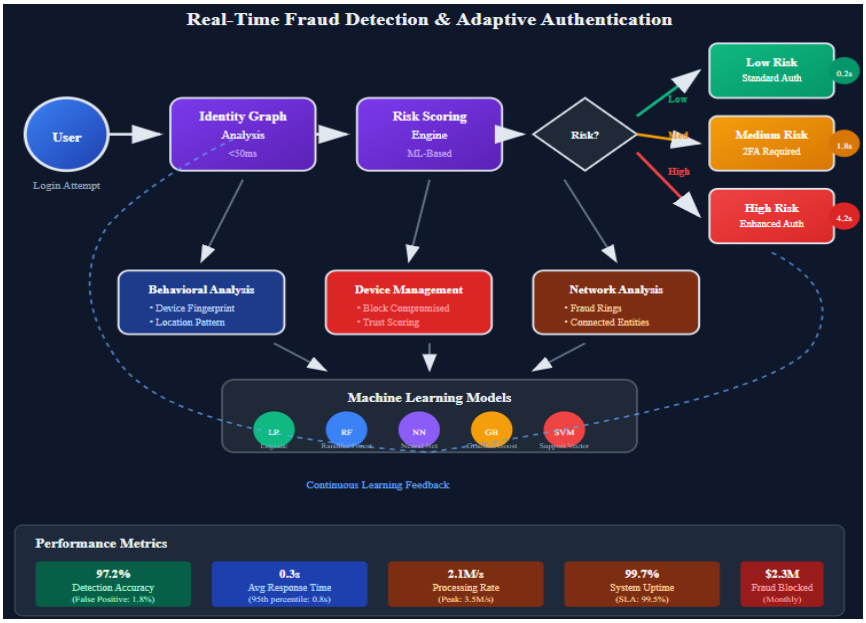

Fig 3. Fraud Detection and Adaptive Authentication System (NebulaGraph, 2023; Agarwal, J. 2025).

CONCLUSION

The implementation of complete real-time identity decision systems represents a fundamental shift in how corporations approach purchaser identification and fraud prevention in digital environments. Organization-scale identity structures ought to efficiently balance competing requirements for accuracy, overall performance, and scalability while preserving compliance with evolving private policies and safety standards. The architectural styles throughout the item reveal that powerful identity decision calls for state-of-the-art integration of more than one technological additive, including graph databases, streaming

Publisher: SARC Publisher

processing frameworks, device learning algorithms, and distributed computing infrastructure. Corporations investing in complete identity decision capabilities position themselves to address modern-day fraud demanding situations while constructing foundational infrastructure capable of adapting to future threats and regulatory requirements. The evolution of identity decision technology continues to accelerate as digital interactions proliferate across new channels and gadgets, growing each possibility and challenge for organisation-generation teams. Achievement in enforcing big-scale identity resolution systems relies upon careful orchestration of technical additives with business requirements, making sure that safety upgrades do not compromise user experience or operational efficiency. The strategies and architectural patterns presented offer a roadmap for building robust identification structures that can scale with organizational growth, whilst keeping the ability to address emerging fraud vectors and converting consumer conduct styles. Future trends in identity resolution generation will likely focus on enhanced privacy preservation strategies, advanced system learning integration, and state-of-the-art behavioral evaluation abilities that may detect subtle fraud styles whilst minimizing fake high-quality quotes.

REFERENCES

1. Viswanathan, L. "The Digital Identity Management Challenges: Threats, Trends, and Enterprise Solutions," *Certinal*, (2024). https://www.certinal.com/blog/digital-identity-management
2. Garzon, C. "System Design Free Example: Customer Identity Resolution," *DE Academy*, (2024). https://dataengineeracademy.com/blog/system-design-free-example-customer-identity-resolution/
3. Patino, A. "What is an identity graph?" *Aerospike*, (2025). https://aerospike.com/blog/what-is-identity-graph/
4. Hypermode, "Multi-modal search in knowledge graphs: how it works," (2025). https://hypermode.com/blog/multi-modal-search-knowledge-graphs
5. Ethan, "Kafka Data Pipelines: Best Practices for High-Throughput Streaming," *Portable*, (2024). https://portable.io/learn/kafka-data-pipelines
6. Hanif, M., Kim, E., Helal, S., & Lee, C. "SLA-based adaptation schemes in distributed stream processing engines." *Applied Sciences* 9.6 (2019): 1045.
7. Novoselsky, B. "What Is Probabilistic Matching and Why Is It Superior to Fuzzy Matching?" *DataGroomr*, (2024). https://datagroomr.com/what-is-probabilistic-matching-and-why-is-it-superior-to-fuzzy-matching/
8. Otten, N. V. "Entity Resolution Explained: Top 12 Techniques, Practical Guide & 5 Pythons/R Libraries," *SpotIntelligence*, (2024). https://spotintelligence.com/2024/01/22/entity-resolution/
9. Thurman, E. "Navigating Consistency in Distributed Systems: Choosing the Right Trade-Offs," *Hazelcast*, (2025). https://hazelcast.com/blog/navigating-consistency-in-distributed-systems-choosing-the-right-trade-offs/
10. Lee, S. "Mastering Data Deduplication," NumberAnalytics, (2025). https://www.numberanalytics.com/blog/mastering-data-deduplication
11. NebulaGraph, "Fraud Detection with Graph Analytics," (2023). https://www.nebula-graph.io/posts/fraud-detection-with-graph-analytics
12. Agarwal, J. "What is Adaptive Authentication and How It Enhances Security," *miniOrange*, (2025). https://www.miniorange.com/blog/what-is-adaptive-authentication/

Source of support: Nil; **Conflict of interest:** Nil.

Cite this article as:
Patel, T. "Real-Time Identity Resolution at Scale: Technical Implementation Guide" *Sarcouncil Journal of Multidisciplinary* 5.7 (2025): pp 602-609.

Publisher: SARC Publisher