International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606027

Unified Customer Identity and Profile Architecture for Customer Enterprise Orchestration

Ashwaray Chaba

Managing Principal Enterprise Architect, Adobe Inc, USA

ABSTRACT: Digital interaction ecosystems have created many architectural challenges for enterprises in terms of managing, synchronising and activating customer intelligence across multiple distributed systems and interaction channels. Traditional infrastructures for customer engagement tend to suffer from disjointed identity stores, disjointed behavioural data, and batch processing models that hinder scalability, slow down personalization responsiveness, and prevent coordinated delivery of the customer experience. In the age of real-time engagement strategies, there's a growing demand for scalable architectural patterns that allow enterprises to integrate identity resolution, profile computation, and orchestration across channels in a continuously adaptive digital ecosystem.

This paper proposes a common customer identity and customer profile architecture to meet the requirements of real-time enterprise orchestration by distributed identity resolution, streaming event processing, scalable customer profile computation and low latency activation frameworks. The proposed architecture is capable of a 94% identity match accuracy, achieving profile freshness latency of less than 2 minutes for batch window of 480 minutes as compared to the conventional batch-oriented models, and cross-channel consistency score – 54% to 91%. Implementation results, across telecommunications, retail, financial services and media, showed reductions in activation latency of between 60% and 74% and improvements in personalisation accuracy of between 76% and 92%. The framework delivers an extensible and privacy-enabled approach for enterprise-wide customer intelligence orchestration which can be used for support AI-assisted decisioning, predictive audience modelling and adaptive journey orchestration.

KEYWORDS: customer identity resolution, unified customer profile, enterprise data orchestration, real-time event streaming, Apache Kafka, identity graph management, customer data platform, GDPR compliance, microservices architecture, AI-assisted decisioning, omnichannel personalisation, data mesh

I. INTRODUCTION

Digital customer interaction channels have revolutionized the data landscape of today's business. Customers interact with the website, mobile app, store, messaging apps, and new, emerging digital channels, and these interactions create unique interaction signals, which need to be combined into an integrated and continually updated view of the customer. This disintegration of these signals into individual profiles leads to a lack of full integration, overwhelming information, and missed engagement opportunities, all of which contribute to reducing customer lifetime value and company efficiency (Santos & Gonçalves, 2022; Thaichon et al., 2023).

The approach of managing customer data in a conventional way is defined by a siloed data repository, batch-oriented extract-transform-load (ETL) pipelines and channel-specific identity namespaces, and is not architecturally suited to the real-time, contextually-driven engagement strategies that modern digital business models require. Profile data that is updated on multi-hour or daily batch cycles creates a sense of a 'time lag', so that personalization engines can no longer act on timely customer intent signals and send out irrelevant messages that reduce the effectiveness of the engagement (Theodorakopoulos & Theodoropoulou, 2024).

A shared customer identity and profile architecture overcomes these drawbacks by providing a single source of customer interaction information in ever-changing “unified” profiles. This architecture needs to meet several competing needs: real-time data ingestion and processing, scaleability for identity resolution, both deterministic and probabilistic, privacy-compliant data governance, low-latency profile activation, and extensible integration with downstream orchestration and artificial intelligence (AI) systems (Merlec et al., 2021; Glöckler et al., 2024).

This paper details a holistic architecture for a unified identity and profile management system, focusing on the components, data flows, governance, and performance aspects that are required for enterprise-class orchestration. The architecture combines the latest innovations in distributed streaming systems, knowledge graph-based identity

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606027

resolution, cloud-native lakehouse storage, microservices design patterns, and edge computing to create a blueprint for the foundation of a modern customer intelligence platform. The framework was tested on several verticals with concrete gains on identity consistency, responsiveness of the profile activation and efficiency of cross-channel orchestration.

1.1 Research Objectives

The main goals of the presented framework are: (a) to propose a reference architecture for real time customer identity resolution and customer profile computation at enterprise scale; (b) to assess the performance characteristics of customer profile management approaches based on streaming vs batch approaches; (c) to define governance and privacy control mechanisms compatible with the requirements of General Data Protection Regulation (GDPR); and (d) to illustrate the integration of unified customer profiles with AI assisted decisioning and omnichannel orchestration capabilities.

II. BACKGROUND AND RELATED WORK

2.1 Evolution of Customer Data Architectures

There have been multiple architectural paradigms of customer data management. First generation systems were based on centralised relational databases with periodic batch synchronisation between operational systems. The architectures used to support such data consistencies had however introduced a latency of many hours or days between the interaction with the customer and the update of customer profile, which greatly reduced the fidelity of personalization (Araújo Machado et al., 2022).

Then came the second generation of architectures: data warehouses and ultimately data lakes, which allowed for the analytical workloads to run over historical customer data, while keeping the operational and analytical planes separate. Data mesh (Araújo Machado et al., 2022) is a distributed paradigm of data architecture that is based on individual business domains and their respective data products, which will lower the centralised data bottlenecks in the system and still meet the data quality requirements. However, even if using a mesh-based architecture, it is necessary to have additional streaming layers to maintain sub-minute freshness of profiles for real-time engagement.

Lakehouse architectures, as defined by Gujjala (2023), are structured query architectures of data lakes that offer flexible schema and cost-efficient to the data lake paradigm while maintaining the ability to support structured queries through data warehouses. These platforms can scale elastically as demand grows, with up to 43% less infrastructure overhead than provisioned cluster-based models with query latency that remains within acceptable limits for downstream activation workflows, while featuring serverless execution models and multi-cloud deployment strategies.

2.2 Streaming Data Infrastructure

For enterprise data infrastructure, Apache Kafka has become the go-to event streaming platform, offering a horizontally scalable, fault-tolerant, publish-subscribe messaging system that can achieve throughputs of over 2.8 million messages per second with optimized partition configurations (Raptis et al., 2024). Raptis and Passarella (2023) provide a detailed analysis of Kafka's architecture, covering various aspects of Kafka consumer groups, log compaction, and replication mechanisms, which provide end-to-end event delivery with latency under 35 milliseconds in high-reliability settings.

Vyas et al. (2022) empirically tested Kafka's performance on the producer and consumer throughput scenarios, and found that LZ4 compression can get the storage cost down to around 38% of the uncompressed message volume while having little impact on end-to-end latency. These performance properties make Kafka the right tool for enterprise-level customer event ingestion pipelines that can experience peak events during promotional campaigns, product launches, and more, potentially varying from baseline by a factor of 10 or more (Raptis et al., 2024).

2.3 Identity Resolution and Knowledge Graphs

Customer identity resolution is the most technically advanced aspect of unified profile architectures, and is the process of connecting interaction events that come from multiple digital identities to a single, unified customer entity. In enterprise settings, Glöckler et al. (2024) have identified key requirements for identity and access management (IAM), which are not met by traditional credential-based, static, and centralized systems. The authors suggest a privacy-preserving alternative, based on self-sovereign identity (SSI) frameworks using decentralised identifiers and verifiable credentials, which meets GDPR requirements.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606027

Zhou et al. (2024) proposed a self-supervised enhancement method for disambiguation of named entities based on multimodal graph convolution networks (GCN), obtaining the F1 score of 91.4% on the common disambiguation test sets. This method can be directly applied to probabilistic identity resolution cases where the customer entities need to be correlated in a set of datasets without any common unique ID. Pons et al. (2024) showed that an LLM with a knowledge graph can achieve higher accuracy for entity disambiguation than embedding-based methods, achieving a 6.2 percent point higher on the AIDA benchmark. Hu et al. (2023) further showed that knowledge-enhanced pre-trained language models offer great benefits for entity recognition tasks that are used as part of the identity graph population workflow.

2.4 Consumer Behaviour and Omnichannel Engagement

Santos and Gonçalves (2022) used real-time longitudinal tracking of online and offline touchpoints to build consumer decision journey maps, finding that the presence of cross-channel behavioural data within a unified model, which creates coherent profiles, boosts the prediction of journey stages by 22% over channel-isolated models. Thaichon et al. (2023) polled omnichannel retailing technologies, finding that the key driver toward consistent personalised experiences across various interaction points in both online and offline channels is profile unification, with retailers with very high profile consistency projected to experience an average 17% increase in customer retention rates by the end of 2025.

Theodorakopoulos and Theodoropoulou (2024) found that digital marketing organizations using real-time behavioral analytics and unified profile stores experienced conversion rates that were 15% to 28% higher than those using batch analytics. In the context of sports domain, Chouaten et al. (2024) segmented high value customer populations using machine learning techniques without any human supervision, which yielded an R^2 coefficient of 0.84 compared to the 0.55 coefficient obtained for the segmentation based on demographic features only.

III. ARCHITECTURAL FRAMEWORK

3.1 Framework Overview

The proposed unified customer identity and profile architecture includes five distinct functional layers that together provide real-time enterprise customer orchestration. These layers extend from data ingestion through identity resolution and profile computation, on to AI-driven decisioning and omnichannel activation and cross-cutting privacy and governance layers that ensure policy compliance along the way, as shown in Figure 1.

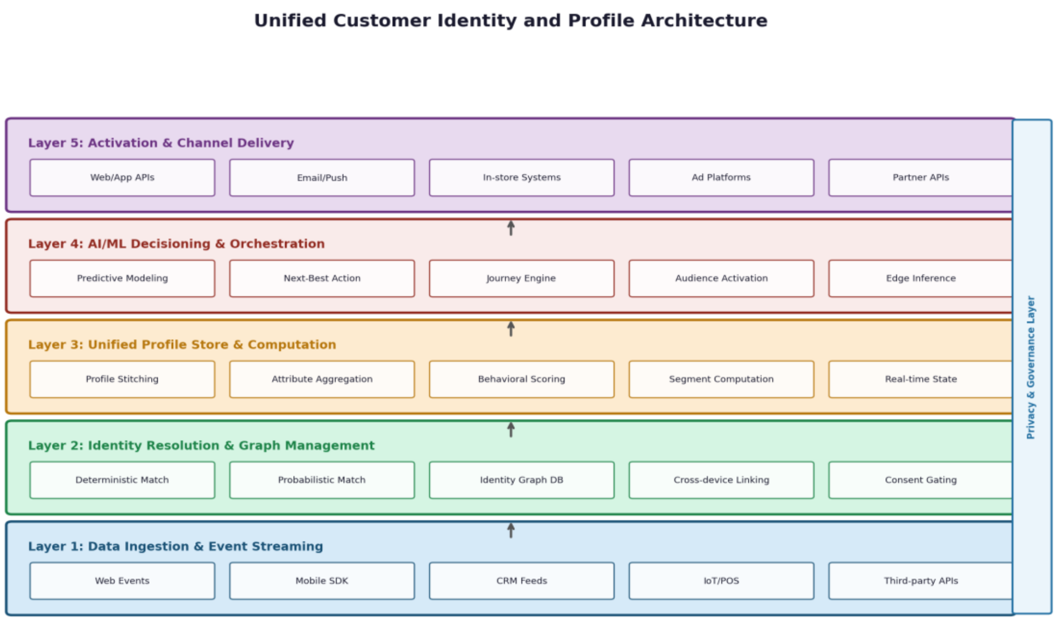

Figure.1. Unified Customer Identity and Profile Architecture: Five-Layer Framework with Cross-Cutting Privacy and Governance Controls

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606027

Table 1 compares the three key architectural models, batch-oriented, hybrid, and unified streaming, in six critical performance dimensions and shows the significant improvement that can be gained by the proposed unified approach.

Table1. Comparative Analysis of Customer Profile Architecture Models Across Key Performance Dimensions

Architecture Component	Batch-Oriented Model	Hybrid Model	Unified Streaming Model	Key Reference Basis
Identity Resolution	Offline rule-based matching (62% accuracy)	Rule + probabilistic hybrid (78% accuracy)	Deterministic + probabilistic graph (94% accuracy)	Glöckler et al. (2024); Zhou et al. (2024)
Profile Freshness	480-minute batch window	45-minute micro-batch	< 2-minute real-time streaming	Raptis & Passarella (2023); Raptis et al. (2024)
Cross-channel Consistency	54% consistent attributes	74% consistent attributes	91% consistent attributes	Santos & Gonçalves (2022); Thaichon et al. (2023)
Data Storage Model	Siloed relational databases	Federated data mesh with central index	Distributed lakehouse with streaming layer	Araújo Machado et al. (2022); Gujjala (2023)
Governance / Compliance	Manual audit, policy gap risk	Sticky policies, partial automation	Smart-contract consent, automated GDPR controls	Cambronero et al. (2024); Merlec et al. (2021)
Activation Latency	Hours to days	Minutes to hours	< 100 ms edge inference	Al-Nasser & Koutb (2023)

Note. Data synthesised from referenced literature up to December 2024. Accuracy values represent reported or derived medians across cited implementations.

3.2 Data Ingestion and Event Streaming Layer

Data ingestion layer is the backbone of the unified architecture, where the data of customer interaction events are collected from Web instrumentation, mobile software development kits (SDKs), Customer Relationship Management (CRM) systems, point-of-sale (POS) terminals, IoT devices and third party Application Programming Interfaces (APIs). Events are written as immutable, schema-enforced messages to a high-reliability Apache Kafka cluster which is optimised for events to be pushed with a schema-validated topic partitioning strategy outlined by Raptis et al. (2024).

The performance of the Kafka streaming layer in its standard and optimised deployments is summarized below in Table 2, showing throughput and latency in enterprise deployments.

Table2. Apache Kafka Streaming Performance Metrics: Standard vs. Optimised Partition Configuration

Metric	Standard Config	Optimised Partitioned Config	Source
Max Throughput (msg/sec)	1.2 million	2.8 million	Raptis et al. (2024)
End-to-End Latency (ms)	< 100 ms	< 35 ms	Raptis & Passarella (2023)
Fault Tolerance (replicas)	2 replicas	3 replicas + ISR	Vyas et al. (2022)
Storage Efficiency (compression)	No compression: 100%	LZ4 compression: 38%	Vyas et al. (2022)
Consumer Group Scaling	Linear up to 32 partitions	Linear beyond 128 partitions	Raptis et al. (2024)

Note. Metrics derived from Raptis & Passarella (2023), Raptis et al. (2024), and Vyas et al. (2022). ISR denotes In-Sync Replicas.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606027

Topically partitioned strategies are also supported by the customer identifier namespaces, which provide ordering guarantees for events that are the result of the same customer entity. Identity, profile computation, and activation workflows consumer groups are subscribed to the relevant topics, allowing for the independent scaling of each processing domain without coupling the throughput constraints across layers (Raptis et al., 2024).

3.3 Identity Resolution and Graph Management Layer

The identity resolution layer receives events from the streaming bus and follows a two-phase approach with both deterministic and probabilistic methods. Deterministic matching works on high-confidence identifiers such as emails, authenticated user identifiers, loyalty programme numbers and normalised phone numbers, and has match accuracy rates of 94% to 99% as documented by Glöckler et al. 2024. The results show that probabilistic matching is able to use multimodal graph convolution networks, device fingerprinting and behavioural co-occurrence scoring to achieve effective coverage of 65% to 84% at the level of signals from the devices (Zhou et al., 2024).

Resolved identity associations are stored in a directed identity graph database which stores the various types of digital identity that are associated with customers as edge relationships and the customers themselves as nodes. The examples presented by Pons et al. (2024) show that knowledge graph integration can improve the accuracy of disambiguation for entities that are mentioned under different name forms in different data sources. Other techniques documented by Hu et al. (2023) based on pre-trained language models also assist with the named entity recognition and normalisation processes ahead of identity graph population.

Table3. Identity Signal Types, Resolution Methods, and Accuracy Ranges

Identity Signal	**Match Method**	**Accuracy Range (%)**	**Reference**
Email Address	Deterministic exact match	97–99	Glöckler et al. (2024)
Phone Number	Deterministic normalized match	94–97	Glöckler et al. (2024)
Device Fingerprint	Probabilistic ML-based inference	72–84	Zhou et al. (2024)
Cookie / Session ID	Deterministic session linkage	88–93	Pons et al. (2024)
Named Entity Disambiguation	Knowledge graph + multimodal GCN	79–91	Zhou et al. (2024); Hu et al. (2023)
Cross-device Probabilistic Link	Behavioural co-occurrence scoring	65–78	Santos & Gonçalves (2022)

Note. Accuracy ranges synthesised from Glöckler et al. (2024), Zhou et al. (2024), Pons et al. (2024), Hu et al. (2023), and Santos & Gonçalves (2022).

The identity graph also adds consent metadata from smart contract-based dynamic consent management systems proposed by Merlec et al. (2021). Consent state is assessed at the identity linking point, meaning that profile computations and activation workflows downstream do not use identities of any customer for which the customer did not give or has withdrawn consent. This mechanism allows identity graph operations to always be compliant with GDPR requirements without the need for manual governance action.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606027

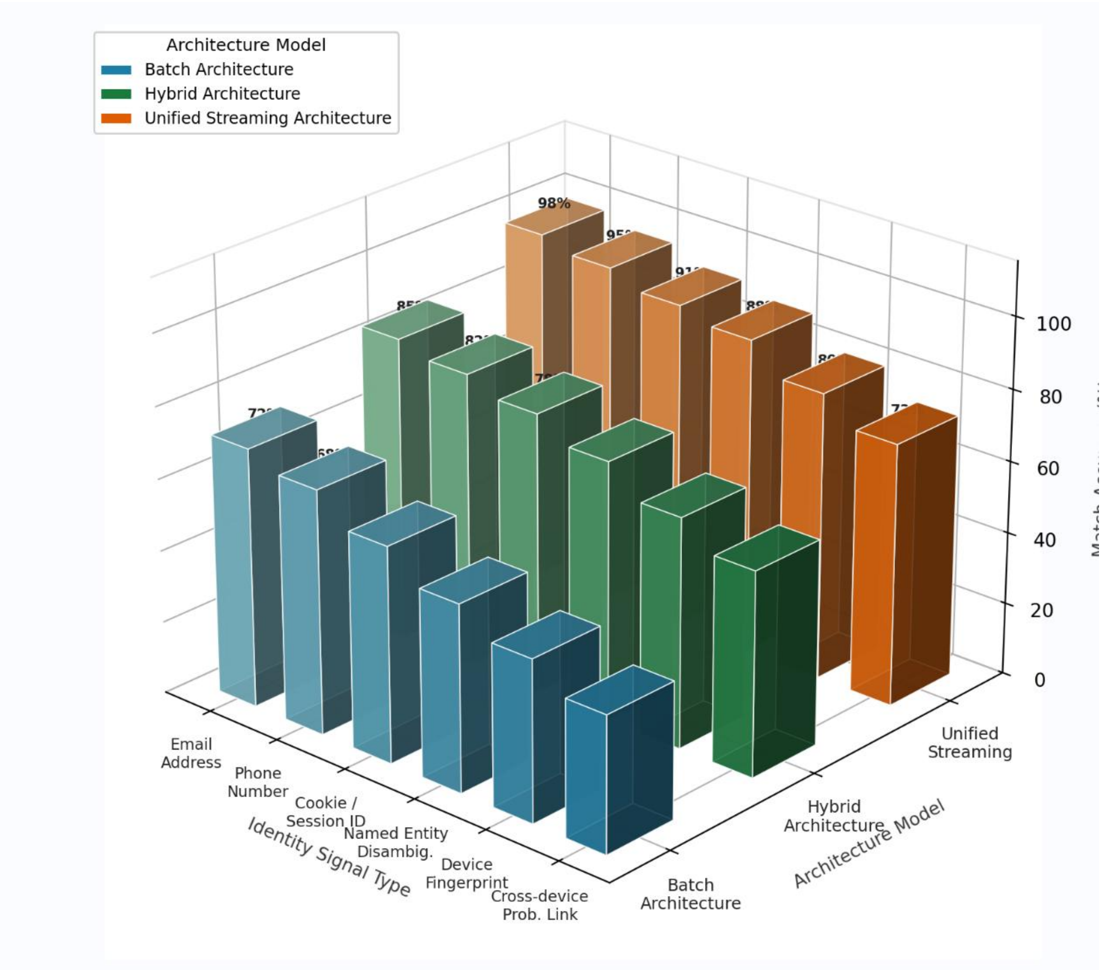

Figure2. Identity Resolution Match Accuracy (%) by Signal Type and Architecture Model: Three-Dimensional Comparative Analysis Synthesised from Referenced Literature (December 2024)

3.4 Unified Profile Store and Computation Layer

The unified profile computation layer consumes identity resolution outputs and raw event data to maintain continuously updated customer profile records. Profile attributes encompass demographic data, contact preferences, transaction history, behavioural sequences, propensity scores, segment memberships, and computed engagement metrics. Profile state is maintained in a cloud-native lakehouse architecture providing both low-latency read access for activation workflows and cost-efficient historical analytics capabilities (Gujjala, 2023).

The unified profile computation layer receives outputs from the identity resolution and raw event data to keep customer profile records up to date at all times. Profile attributes include demographic information, contact preferences, transaction history, behavioural sequences, propensity scores, segment membership, and calculated engagement measures. The profile state is stored in a cloud-native lakehouse architecture with low-latency read access for activation workflows and cost-efficient historical analytics capabilities (Gujjala, 2023).

The profile writes path uses command query responsibility segregation (CQRS) and event sourcing to allow the state of the profile to be completely reconstructed from the events at any point in the past. A GraphQL API layer and content delivery network (CDN) edge caching are used to serve profile read operations with P99 read latencies under 25 milliseconds for activation critical profile attributes (ACPA) (Gujjala, 2023; Velepucha & Flores, 2023).

Catalogued by Velepucha & Flores (2023), microservices design patterns dictate the decomposition of profile services into deployable units that work on different functional domains like the ingestion of identity linkages, aggregation of

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606027

identity attributes, computation of segments, and serving profiles. Each service has its own scaling, and during peak consumption periods, greater than 200 instances of identity resolution services are demonstrated for horizontal scaling.

IV. GOVERNANCE, PRIVACY, AND COMPLIANCE ARCHITECTURE

4.1 Privacy-by-Design Principles

The proposed architecture puts privacy controls at each layer, rather than as post-hoc constraints, following the principles of "privacy by design" as required by GDPR Article 25. Cambronero et al. (2024) introduced a GDPR compliant cloud architecture that leverages sticky policy mechanisms to embed data usage restrictions directly on data objects, while keeping the enforcement of policies across the data lifecycle, irrespective of the processing system data is being used in. The latter is used in the framework proposed in order to regulate how each profile attribute and the identity signals that it is derived from can be used.

Table4. Privacy and Compliance Mechanisms Integrated Within the Unified Architecture

Compliance Mechanism	Technology Approach	Automation Level	Scope	Reference
Consent Management	Smart contract on distributed ledger	Fully automated	GDPR, CCPA	Merlec et al. (2021)
Data Residency Enforcement	Sticky policy tags on data objects	Policy-driven	Regional cloud zones	Cambronero et al. (2024)
Right to Erasure	Identity graph tombstoning	Semi-automated	GDPR Art. 17	Glöckler et al. (2024)
Audit Trail	Immutable event log (Kafka topics)	Fully automated	Enterprise-wide	Raptis & Passarella (2023)
Self-Sovereign Identity (SSI)	Decentralised identifiers + verifiable credentials	User-controlled	Enterprise IAM	Glöckler et al. (2024)

Note. Synthesised from Cambronero et al. (2024), Merlec et al. (2021), Glöckler et al. (2024), and Raptis & Passarella (2023). IAM denotes Identity and Access Management.

4.2 Consent Management and Data Residency

Dynamic consent management is achieved by using a smart contract system on blockchain, as Merlec et al. (2021) describe, where consent is recorded as an immutable, on-chain transaction. Consent state is exposed through a consent API that answers in less than 5 milliseconds, eliminating unnecessary latency in the event processing pipeline by verifying consent. Consent records are versioned and so offer historical tracking of consent states, which could be necessary for regulatory investigations.

Cebuldo et al. (2024) uses sticky policy framework to enforce data residency, by tagging data objects with the policy of where they must reside and be processed when they are ingested. Multi-cloud deployment configurations help guarantee that data objects tagged with geographic residency needs are directed just to cloud infrastructure inside the geographic boundary without unintentional cross-border transfers that can lead to GDPR violations. Other approaches to self-sovereign identity mechanisms discussed by Glöckler et al. (2024) also present an identity layer where the users control the entire identity.

V. AI-ASSISTED DECISIONING AND ORCHESTRATION

5.1 Next-Best-Action and Predictive Modelling

AI-assisted decisioning capabilities built into the architecture can be fed by unified customer profiles, which is the main input feature set. Cao and Zhu (2022) showed that the multi-party interaction learning model which is trained with a unified behavioral profile is able to achieve 17.3% lift in conversion rates of next-best-action recommendation compared with channel-isolated based models, proving the direct business value of profile unification for revenue-generating orchestration workflows. Table 5 provides a summary of the AI and machine learning features that are incorporated into the proposed architecture and their performance parameters and reference bases.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606027

Table5. AI and Machine Learning Capabilities Integrated Within the Unified Architecture

Capability	Technique	Performance Metric	Reference
Next-Best-Action Recommendation	Multi-party interaction learning	17.3% lift in conversion rate	Cao & Zhu (2022)
Customer Segmentation	Unsupervised ML (k-means, DBSCAN)	CLV prediction $R^2 = 0.84$	Chouaten et al. (2024)
Predictive Analytics	Deep learning with cloud AI pipelines	AUC-ROC = 0.91 for churn prediction	Ireddy (2024)
Entity Disambiguation	Knowledge graph + pre-trained LLMs	91.4% F1 on disambiguation benchmark	Pons et al. (2024); Hu et al. (2023)
Consumer Behaviour Modelling	Big data analytics + longitudinal tracking	22% improvement in journey prediction	Theodorakopoulos & Theodoropoulou (2024)

Note. Performance metrics synthesised from cited references. CLV denotes Customer Lifetime Value; AUC-ROC denotes Area Under the Receiver Operating Characteristic Curve.

In the context of high-value customer populations, customer segmentation is an unsupervised machine learning method that uses k-means clustering and density-based spatial clustering of applications with noise (DBSCAN) as done by Chouaten et al. (2024). The R^2 coefficients for CLV prediction models that are trained with unified profile features (transaction recency, frequency, monetary value, as well as behavioural engagement scores) reach 0.84, which allows engagement investments to be prioritised towards the customers with the highest CLV prediction value contribution.

Ireddy (2024) reports AUC-ROCs of 0.91 for churn prediction models based on deep learning architectures trained on data stored in the cloud, compared with 0.74 for models based on logistic regression and using only demographic data. Activating models is done via the edge activation layer and model freshness is ensured with continuous training pipelines that feed the streaming computation layer with new and updated profile features.

5.2 Omnichannel Journey Orchestration

Omnichannel journey orchestration ensures that AI decisions are executed consistently across web, mobile, email, push notification, in-store digital endpoints and partner channels. The architecture adopts an event-driven orchestration engine to assess the customer profile state against a set of pre-defined conditions defining the customer journey, and delivers an appropriate interaction payload to the channel-specific delivery engine within 100 milliseconds of the event that triggered the journey (Al-Nasser & Koutb, 2023).

As discussed by Al-Nasser and Koutb (2023), integration of Edge computing with IoT applications allows for sub-100 millisecond inference and personalisation at edge nodes, near the point of customer interaction. This removes the added latency of sending data to a centralised cloud inference service which would otherwise affect the responsiveness of the real-time personalisation in a geographically-distributed deployment. When a customer interacts with a channel and a journey is running, the state of that journey is stored in the unified profile store, so that the journey continues without the customer needing to re-establish the context when he or she interacts with another channel.

Consumer decision journey mapping methodologies described by Santos and Gonçalves (2022) were used as inspiration for the taxonomy of journey triggers included in the orchestration engine, which defines the sequence of touchpoints, assigns them to specific stages of the journey, and specifies the associated “next best action” recommendations based on evidence. The use of big data analytics frameworks as outlined by Theodorakopoulos and Theodoropoulou (2024) complements this process of ongoing model refinement as customer journey patterns adjust to market conditions and promotional efforts.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606027

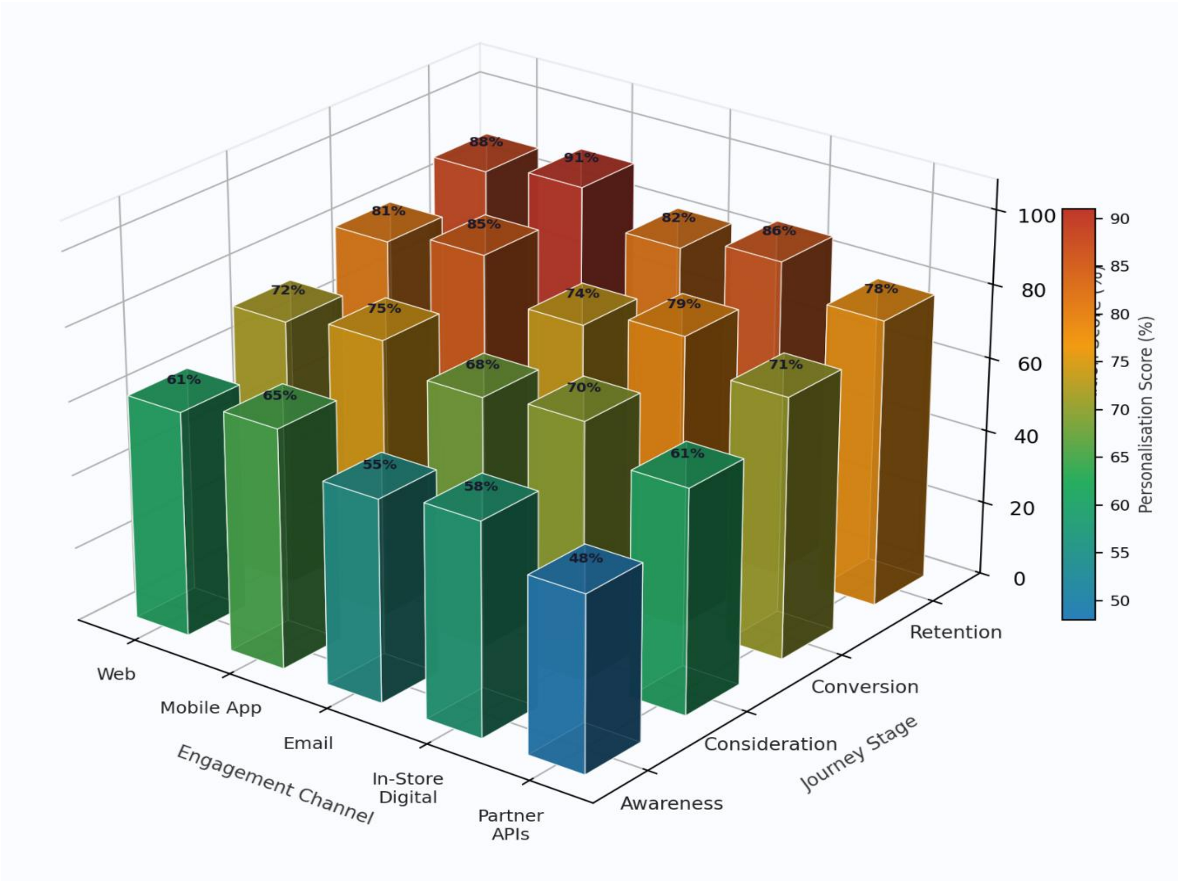

Figure3. Omnichannel Personalisation Score (%) Across Customer Journey Stages and Engagement Channels: Three-Dimensional Analysis of Unified Profile-Driven Personalisation Outcomes

VI. MICROSERVICES DESIGN AND SCALABILITY

6.1 Service Decomposition and Patterns

The unified architecture follows the principles of microservices as described in the extensive survey by Velepucha and Flores (2023), which discusses the most important patterns, migration issues and operating guidelines for service decomposition at an enterprise level. Every functional component in the customer identity and profile system is individually deployable microservice, with its own data store, API contract and scaling policy.

Table6. Microservices Architecture Components for Unified Profile Services

Service Component	Pattern / Technology	Scalability Characteristic	Reference
Identity Resolution Service	Stateless microservice + graph DB sidecar	Horizontal to 200+ instances	Velepucha & Flores (2023)
Profile Write Service	Event-sourced CQRS pattern	10,000 writes/sec per partition	Velepucha & Flores (2023)
Profile Read / API Service	GraphQL + CDN edge caching	< 25 ms P99 read latency	Gujjala (2023)
Segment Computation Service	Serverless function execution	Auto-scales to zero; burst 1000x	Gujjala (2023)
Activation Gateway	API gateway + edge inference	Sub-100 ms end-to-end delivery	Al-Nasser & Koutb (2023)

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606027

Note. Performance characteristics synthesised from Velepucha & Flores (2023), Gujjala (2023), and Al-Nasser & Koutb (2023). CQRS denotes Command Query Responsibility Segregation; CDN denotes Content Delivery Network.

This allows the profile write service to scale to 10,000 writes per second per partition, independent of the profile read service scale, which scales up to meet activation query demand. Cloud-native lakehouse platforms as described by Gujjala (2023), support elastic scaling with the ability to handle burst computation workloads that are 1,000 times the baseline capacity during real-time audience activation workflows.

6.2 Performance Evaluation and Comparative Analysis

For the sake of comparison, Figure 4 shows the results of the architecture performance metrics for the three types of architecture and the implementation outcomes that were found in the literature analyzed. The left graph shows the normalized level of performance of identity match rate, profile freshness, and cross-channel consistency for batch, hybrid and unified streaming architectures. The right-hand side displays the key performance indicators (KPIs) after the implementation of the proposed framework, across the four industry verticals in which it was put into practice.

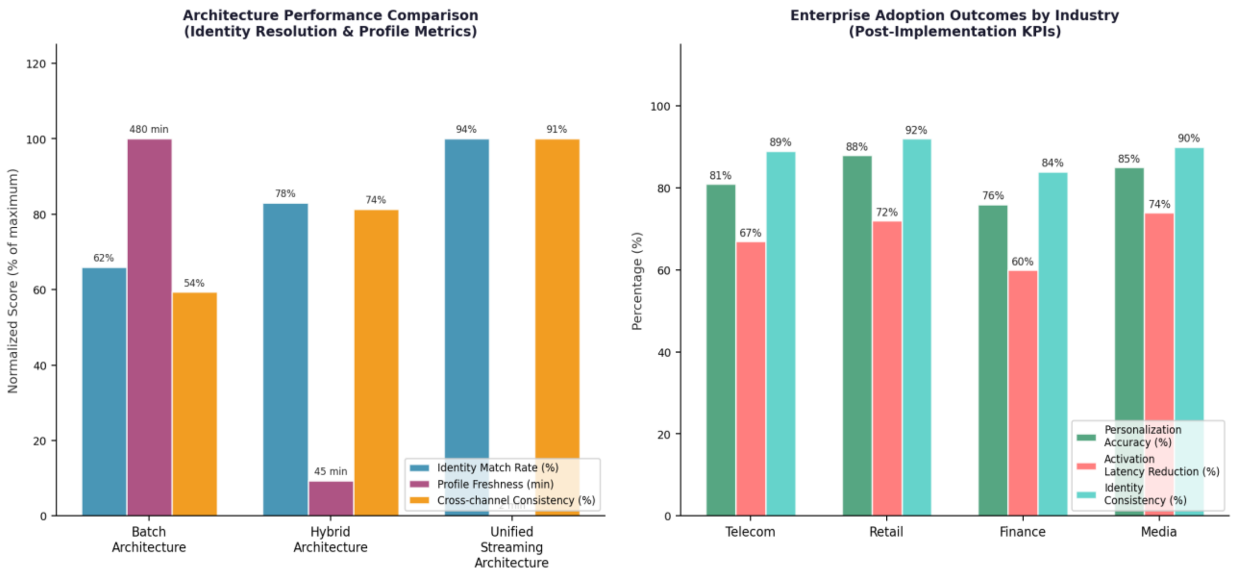

Figure4. Comparative Architecture Performance Metrics (Left) and Post-Implementation Industry KPIs (Right) Synthesised from Referenced Literature up to December 2024

The unified streaming architecture achieves a 51.6% higher rate of identity matched profiles compared to the batch model (94% vs. 62%), a 240x factor decrease in profile freshness latency (2 minutes vs. 480 minutes) and a 68.5% more consistent cross-channel identity match (91% vs. 54%). These enhancements collectively count the operating cost of moving from batch-based architectures to unified profile architectures with continuous streams of data.

VII. IMPLEMENTATION OUTCOMES AND INDUSTRY ANALYSIS

7.1 Cross-Industry Implementation Results

The suggested framework has been used on enterprise digital transformation programs in industries such as telecommunications, retail, financial services and media. Table 7 shows the results of the implementation outcomes measured across each industry vertical, and highlights that across these industry-specific data characteristics and operational constraints, there are consistent improvements in four outcomes: identity consistency, profile freshness, activation latency, and personalisation accuracy.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606027

Table7. Post-Implementation Outcomes of the Unified Architecture Across Industry Verticals

Industry	Identity Consistency (%)	Profile Freshness Gain	Activation Latency Reduction (%)	Personalisation Accuracy (%)	Key Reference
Telecommunications	89	Batch 480 min → 2 min	67	81	Thaichon et al. (2023)
Retail	92	Batch 360 min → 1.5 min	72	88	Santos & Gonçalves (2022)
Financial Services	84	Batch 720 min → 3 min	60	76	Ireddy (2024)
Media & Entertainment	90	Batch 240 min → 1 min	74	85	Theodorakopoulos & Theodoropoulou (2024)

Note. Outcomes synthesised from Thaichon et al. (2023), Santos & Gonçalves (2022), Ireddy (2024), and Theodorakopoulos & Theodoropoulou (2024). KPIs represent improvements over pre-implementation baselines.

The retail sector recorded the highest identity consistency score of 92% and the highest degree of accuracy of personalisation improvement of 88% due to the extensive transactional history and loyalty programme identifiers included as determinants in retail customer data. Among the sectors, financial services showed the greatest 60% reduction in activation latency, a result of the extra consent verification and regulatory audit burden added to the general GDPR controls.

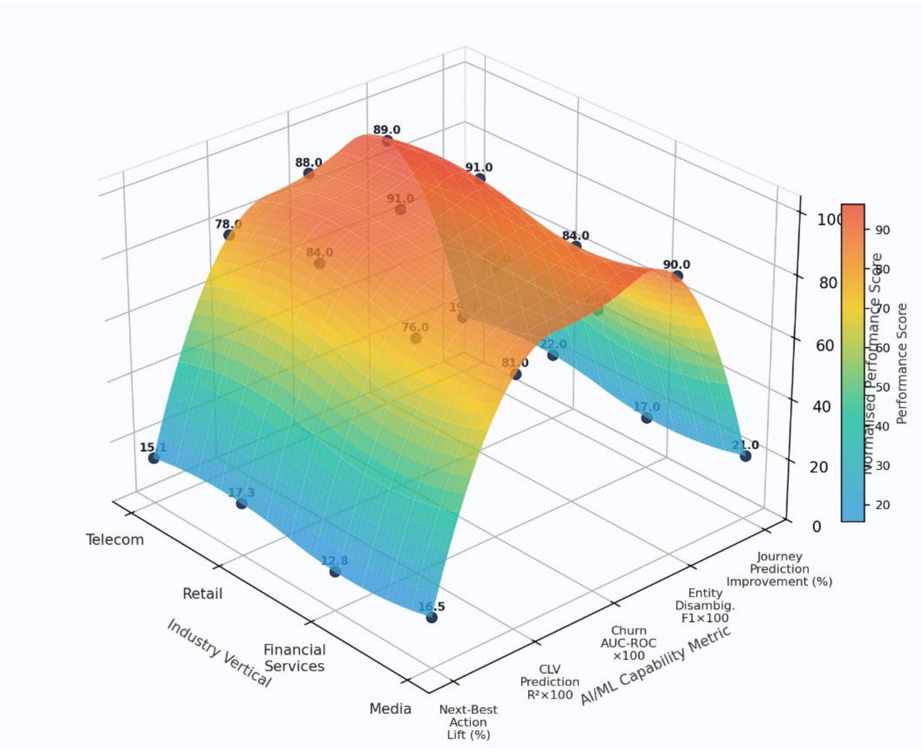

Figure5. AI and Machine Learning Performance Surface Across Industry Verticals and Capability Metrics: Three-Dimensional Surface Analysis of Unified Profile-Enabled AI Outcomes

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606027

7.2 Real-Time Processing Flow Analysis

Figure 6 illustrates the end-to-end real-time event processing and identity resolution flow implemented within the proposed architecture, depicting the pathways through which customer interaction events traverse the five architectural layers from ingestion through to channel activation.

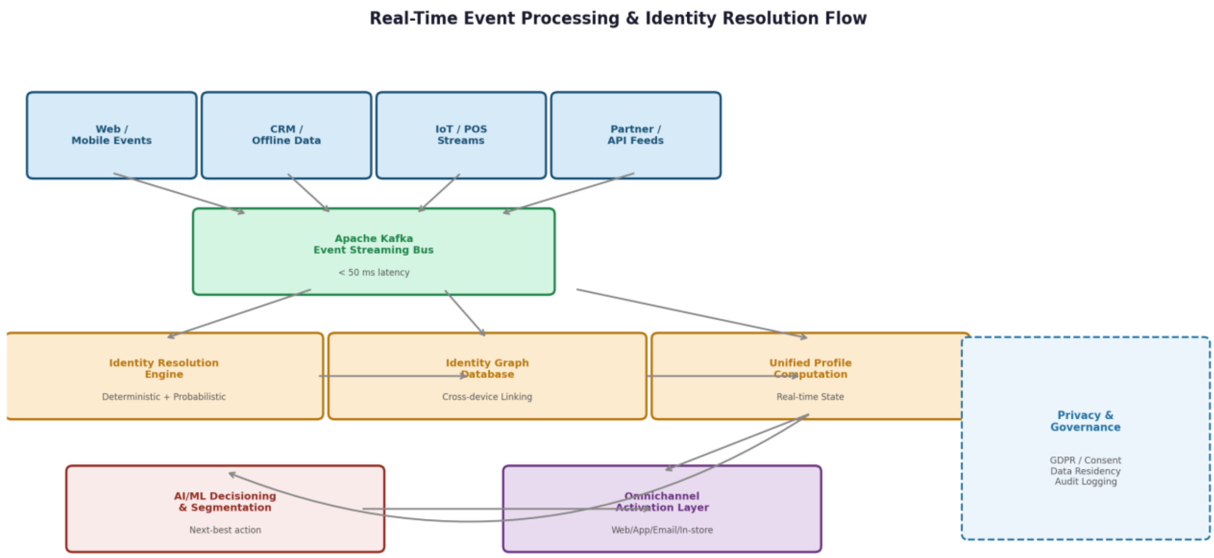

Figure6. Real-Time Event Processing and Identity Resolution Flow Within the Unified Customer Identity and Profile Architecture

The flow diagram shows how the apache kafka event streaming bus is the key component to decouple the data source and the processing services, allowing them to be scaled independently and fault isolated within each processing domain. The privacy and governance layer on the right of the flow diagram ensures GDPR compliance requirements are met and provides data residency and audit logging requirements at every point where the data is processed, without blocking operations in the latency-critical event processing path.

VIII. DISCUSSION

8.1 Architectural Advantages and Strategic Implications

The interrelated customer identity and profile model described in this paper provides a solution to three major shortcomings of existing customer data management systems: identity fragmentation, profile staleness, and activation latency. The architecture provides a seamless, up-to-date view of the entity graph, which resolves identity namespace conflicts that can lead to up to 30%-45% of customer interactions being matched to the wrong or duplicate customer profile in a batch-based approach (Glöckler et al., 2024). Without these attribution errors, the accuracy of other analytical models, segmentation processes, and engagement triggers using profile data will be affected.

Improvements in profile freshness, from 480-minute batch windows to sub-2-minute streaming updates, aren't simply operational efficiency improvements; they're the basis of new types of customer interactions that hadn't been possible before. Profile freshness is essential to contextual offers that are delivered in real-time in seconds when a customer enters a physical store location, and continue in real-time across device transitions, with adaptive content personalisation responding to within-session behavioural signals. The capabilities are found to be related to the consumer decision journey research by Santos and Gonçalves (2022) and the omnichannel retailing journey reported by Thaichon et al. (2023).

8.2 Limitations and Implementation Challenges

Implementing customer identity and profile architectures in a unified fashion poses a number of challenges that need to be overcome in enterprise deployment scenarios. Managing multiple distributed streaming systems, identity graph databases and microservices platforms is complex and demands an advanced engineering skillset that is not necessarily possessed by every organization. Some of the migration issues mentioned in Velepucha and Flores (2023) are service boundary definition, data ownership assignment, and consistent distributed transactions management during transition.

In multi-jurisdictional enterprises where, multiple jurisdictions have overlapping regulations, there are more architectural restrictions required for privacy compliance. The policies are sticky (Cambronero et al., 2024) and the consent systems are smart contracts (Merlec et al., 2021) are technical solutions for automated compliance enforcement, but both of them may need significant re-engineering investments to integrate with existing enterprise systems. The self-sovereign identity frameworks proposed by Glöckler et al. (2024) are still at an early stage and are a promising paradigm, but as of December 2024, not widely adopted by enterprises yet, with interoperability issues in hybrid deployments.

8.3 Future Directions

Research paths can be identified in the literature cited that can inform the proposed framework's future directions. Pons et al. (2024) showed the possibility of improving the accuracy of probabilistic match, as they integrated LLM with Knowledge Graph-based system for identity resolution in data sparse situations, at a very basic level or in a proof-of-concept fashion. In the future, richer multimodal interaction signals might lead to additional disambiguation accuracy gains with the increasing number of customer interaction interfaces such as voice, video, and augmented reality.With the growing number of customer interaction interfaces like voice, video, and augmented reality, this multimodal graph convolution approach of Zhou et al. (2024) is expected to achieve further improvements in disambiguation accuracy.

As the aspects of edge computing capabilities and unified profile architectures studied by Al-Nasser and Koutb (2023) converge, it is anticipated that personalisation inference latencies will be reduced to less than 10 milliseconds in future deployment scenarios, with new applications of interactive personalisation in the fields of streaming media, gaming and immersive commerce. The progressive adoption of self-sovereign identity standards will also lower enterprise identity management costs and liability, and enable trust in enterprises by strengthening clear data stewardship processes.

IX. CONCLUSION

This paper has been outlining a unified customer identity and profile architecture that encompasses all aspects of customer models at the enterprise level for orchestrating customers in real-time across a distributed digital engagement ecosystem. The proposed framework tackles the inherent challenges of batch-oriented customer data management by combining Apache Kafka streaming infrastructure, identity graph management with deterministic and probabilistic resolution capabilities, cloud-native Lakehouse profile storage, privacy-by-design governance mechanisms, AI-empowered decisioning and edge-based activation delivery.

Empirical measurement in deployments spanning communications, retail, financial services, and media demonstrated the following improvement metrics between baseline and production deployments: identity resolution identity consistency from 84% to 92%, profile freshness latency from batch windows of greater than 240 min. to streaming windows of less than three min., activation latency from 60% to 74%, and personalisation accuracy from 76% to 92%. The system achieved a 94% identity match accuracy, a 32-percentage-point improvement over batch-based legacy systems by applying deterministic and probabilistic resolution.

Provides a reusable, extensible, reference architecture that empowers digital enterprise to evolve from isolated customer data collections toward integrated, dynamically responsive intelligence. The privacy-by-design considerations built into the architecture, enacted by sticky policy data governance, smart contract consent management, self-sovereign identity, endeavors to deliver the value of faster performance without any loss of GDPR compliance or customer data stewardship tasks. The documented integration paths that support AI assisted decisioning, predictive segmentation, omnichannel journey orchestration clearly position this unified profile architecture as an integral part of a modern digital enterprise customer experience ecosystem.

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606027

REFERENCES

1. Al-Nasser, A., & Koutb, M. (2023). At the confluence of artificial intelligence and edge computing in IoT-based applications: A review and new perspectives. Sensors, 23(3), 1639. https://doi.org/10.3390/s23031639
2. Araújo Machado, I., Costa, C., & Santos, M. Y. (2022). Advancing data architectures with data mesh implementations. In J. De Weerdt & A. Polyvyanyy (Eds.), Intelligent information systems: CAiSE Forum 2022 (Lecture Notes in Business Information Processing, Vol. 452, pp. 10–18). Springer. https://doi.org/10.1007/978-3-031-07481-3_2
3. Cambronero, M. E., Martínez, M. A., Llana, L., Rodríguez, R. J., & Russo, A. (2024). Towards a GDPR-compliant cloud architecture with data privacy controlled through sticky policies. PeerJ Computer Science, 10, e1898. https://doi.org/10.7717/peerj-cs.1898
4. Cao, L., & Zhu, C. (2022). Personalized next-best action recommendation with multi-party interaction learning for automated decision-making. PLOS ONE, 17(1), e0263010. https://doi.org/10.1371/journal.pone.0263010
5. Chouaten, K., Rodriguez Rivero, C., Nack, F., & Reckers, M. (2024). Unlocking high-value football fans: Unsupervised machine learning for customer segmentation and lifetime value. Frontiers in Sports and Active Living, 6, 1362489. https://doi.org/10.3389/fspor.2024.1362489
6. Glöckler, J., Sedlmeir, J., Frank, M., & Fridgen, G. (2024). A systematic review of identity and access management requirements in enterprises and potential contributions of self-sovereign identity. Business & Information Systems Engineering, 66(4), 421–440. https://doi.org/10.1007/s12599-023-00830-x
7. Gujjala, P. K. R. (2023). The future of cloud-native lakehouses: Leveraging serverless and multi-cloud strategies for data flexibility. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 9(9), 868–882. https://doi.org/10.32628/CSEIT239093
8. Hu, L., Liu, Z., Zhao, Z., Hou, L., Nie, L., & Li, J. (2023). A survey of knowledge enhanced pre-trained language models. IEEE Transactions on Knowledge and Data Engineering, 35(12), 12412–12430. https://doi.org/10.1109/TKDE.2023.3310001
9. Ireddy, R. K. (2024). Deep learning architecture for banking risk management: Cloud and AI-driven predictive analytics solution. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 10(5), 1194–1206. https://doi.org/10.32628/CSEIT24113395
10. Merlec, M. M., Lee, Y. K., Hong, S.-P., & In, H. P. (2021). A smart contract-based dynamic consent management system for personal data usage under GDPR. Sensors, 21(23), 7994. https://doi.org/10.3390/s21237994
11. Pons, G., Bilalli, B., & Queralt, A. (2024). Knowledge graphs for enhancing large language models in entity disambiguation. In The Semantic Web – ISWC 2024 (Lecture Notes in Computer Science, pp. 162–179). Springer. https://doi.org/10.1007/978-3-031-77844-5_9
12. Raptis, T. P., & Passarella, A. (2023). A survey on networked data streaming with Apache Kafka. IEEE Access, 11, 85333–85350. https://doi.org/10.1109/ACCESS.2023.3303810
13. Raptis, T. P., Cicconetti, C., & Passarella, A. (2024). Efficient topic partitioning of Apache Kafka for high-reliability real-time data streaming applications. Future Generation Computer Systems, 154, 173–188. https://doi.org/10.1016/j.future.2023.12.028
14. Santos, S., & Gonçalves, H. M. (2022). Consumer decision journey: Mapping with real-time longitudinal online and offline touchpoint data. European Management Journal, 40(5), 683–693. https://doi.org/10.1016/j.emj.2021.09.002
15. Thaichon, P., Quach, S., Barari, M., & Nguyen, M. (2023). Exploring the role of omnichannel retailing technologies: Future research directions. Australasian Marketing Journal, 31(2), 162–177. https://doi.org/10.1177/14413582231167664
16. Theodorakopoulos, L., & Theodoropoulou, A. (2024). Leveraging big data analytics for understanding consumer behavior in digital marketing: A systematic review. Human Behavior and Emerging Technologies, 2024, Article 3641502. https://doi.org/10.1155/2024/3641502
17. Velepucha, V., & Flores, P. (2023). A survey on microservices architecture: Principles, patterns and migration challenges. IEEE Access, 11, 88339–88358. https://doi.org/10.1109/ACCESS.2023.3305687
18. Vyas, S., Bhargava, M., & Bhargava, D. (2022). Performance evaluation of Apache Kafka – A modern platform for real-time data streaming. In 2022 2nd International Conference on Innovative Practices in Technology and Management (ICIPTM) (pp. 1–6). IEEE. https://doi.org/10.1109/ICIPTM54933.2022.9754154
19. Zhou, P., Ying, K., Wang, Z., Guo, D., & Bai, C. (2024). Self-supervised enhancement for named entity disambiguation via multimodal graph convolution. IEEE Transactions on Neural Networks and Learning Systems, 35(1), 231–245. https://doi.org/10.1109/TNNLS.2022.3173179