



Real-Time Entity-Level Identity Resolution

Achieving a True 360° Customer View

A comprehensive guide for data engineers, data scientists, product managers, and business leaders navigating the complexities of modern customer data strategy.

EBOOK



Introduction

In today's fragmented digital ecosystem, customers interact with businesses across a dizzying array of devices, platforms, and channels. Each of these interactions creates a distinct data point, but without a unifying layer, these points remain isolated.



The Data Gap

Businesses collect more data than ever, yet insight remains shallow due to architectural failures.

CRMs, warehouses, and marketing platforms were not designed to speak to one another, leaving customer identifiers inconsistent.



What Identity Resolution Does

It bridges the gap by connecting fragmented signals across systems, channels, and time. It builds a unified and actionable profile of a real-world entity, transforming disconnected fragments into a living identity.



What This eBook Covers

A practical guide moving from first principles to implementation. We examine the technical mechanisms, privacy obligations, organisational considerations, and measurable outcomes needed to succeed.

The Identity Problem in Modern Businesses

At the heart of data strategy failures lies a simple problem: the same person appears differently across systems. A customer using a personal email, a work account, and different devices is seen by most enterprise architectures as multiple separate people.

The Fragmentation Problem

Channels create distinct identifiers without native reconciliation.

- Mismatched emails and device IDs
- Inconsistent CRM naming conventions
- Offline vs. online disconnect
- Schema conflicts in third-party data

The Business Consequences

- **Marketing:** Wasted spend and duplicate outreach
- **Personalization:** Irrelevant recommendations
- **Analytics:** Inaccurate cohort and retention metrics
- **Experience:** Redundant verification steps
- **Fraud:** Weak cross-account detection

30%

Duplicate Records

Portion of records that are fragments of a single identity.

6.6

Devices Per User

Average connected devices per consumer.

\$15M

Annual Waste

Estimated marketing spend lost on misidentified audiences.

What is Real-Time Entity-Level Identity Resolution?

Real-Time Entity-Level Identity Resolution is the process of linking records across data sources into a single, authoritative representation of a real-world entity. This "entity" can represent an individual, household, business account, or device meaningful to your organisation.

Deduplication

- Focuses on exact matches within a single system.
- Merges rows sharing identical identifiers like email addresses.

Identity Resolution

- Handles ambiguity across disparate structured and unstructured sources.
- Reconciles records that share no common identifiers.

Record Identification

The process parses and cleans raw data to extract candidate identifiers across systems.

Record Linkage

Matching logic evaluates candidate pairs to determine the probability they represent the same entity.

Unified Profile Creation

Confirmed matches are merged into an authoritative profile using conflict resolution rules.

Continuous Refinement

The system constantly ingests new signals to update linkages as entities evolve over time.

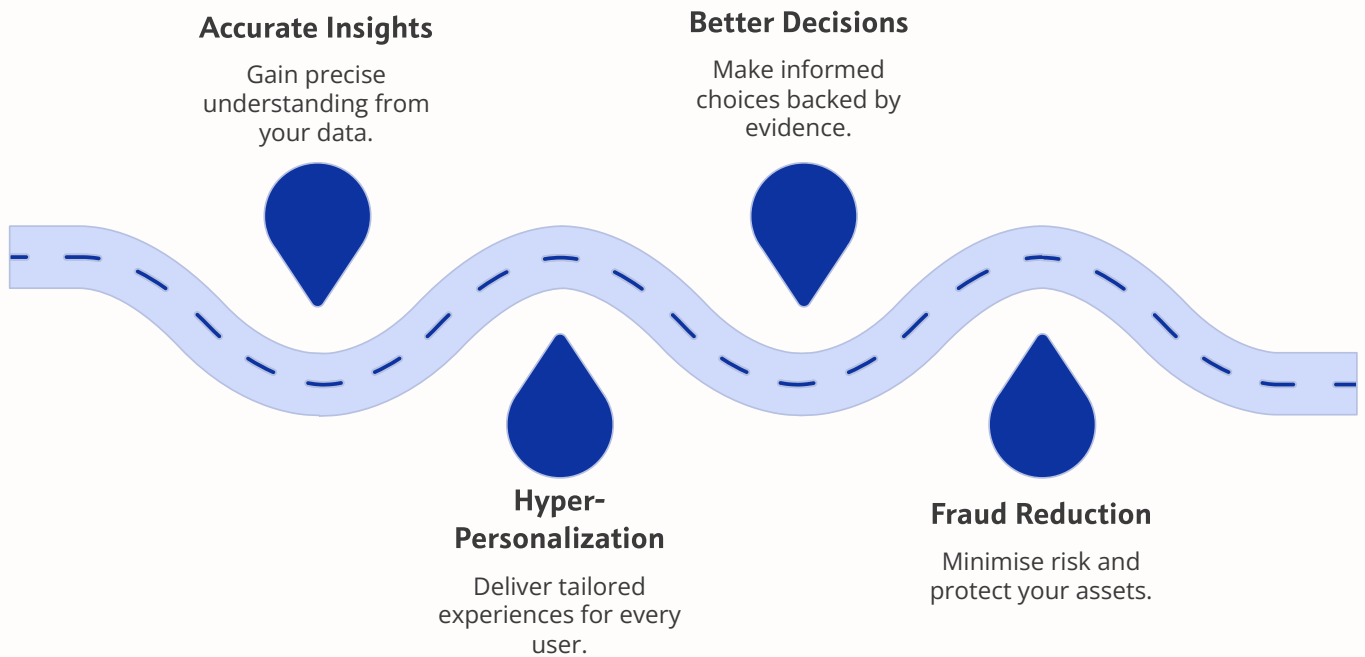
Identity resolution operates at three maturity levels: **Basic** (resolving within a single system like a CRM), **Intermediate** (linking across different systems like web analytics), and **Advanced** (resolving across time and changing contexts).



Why a 360° Customer View Matters


The phrase "360° customer view" has become something of a marketing cliché, but the underlying concept represents a genuine and significant shift in what is possible when data is properly unified. A true 360° view is not simply a comprehensive profile page in a CRM — it is a dynamic, continuously updated representation of a customer's full relationship with your organisation, drawing on every available signal from every touchpoint, channel, and time period.

What a 360° View Enables



Measurable Business Impact

The business case for identity resolution is well-evidenced across industries. Organisations that have invested in unified identity capabilities consistently report:



Beyond the quantitative metrics, a 360° view creates qualitative competitive advantages that are difficult to replicate. Organisations that truly know their customers — not as a collection of transactions, but as coherent human beings with preferences, histories, and evolving needs — are able to build relationships of genuine value. This translates into brand loyalty, advocacy, and long-term revenue sustainability that no amount of advertising spend can manufacture.

How It Works??

Core Components of Identity Resolution

A production-grade identity resolution system relies on five interdependent components that must work in concert to build reliable customer intelligence.



1. Data Ingestion

This layer reliably collects and streams diverse data from CRM, web, and point-of-sale systems.



2. Data Normalisation

The normalisation layer cleans and parses raw inputs into a canonical format to ensure high-quality matching.



3. Matching Engine

This core engine uses rules and models to evaluate and merge records based on their likelihood of belonging to the same individual.



4. Identity Graph

The identity graph functions as a central network of connected identifiers and real-world entities for downstream use.

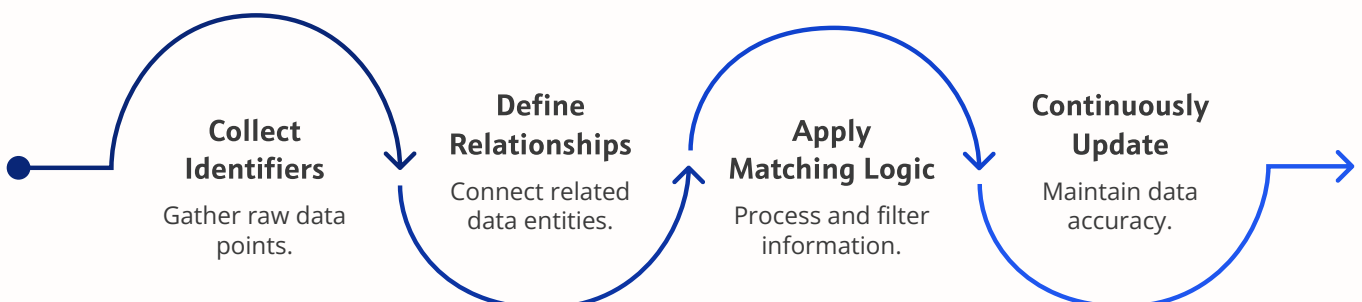


5. Activation Layer

This layer provides secure, real-time access to resolved profiles for marketing, fraud detection, and analytics systems.

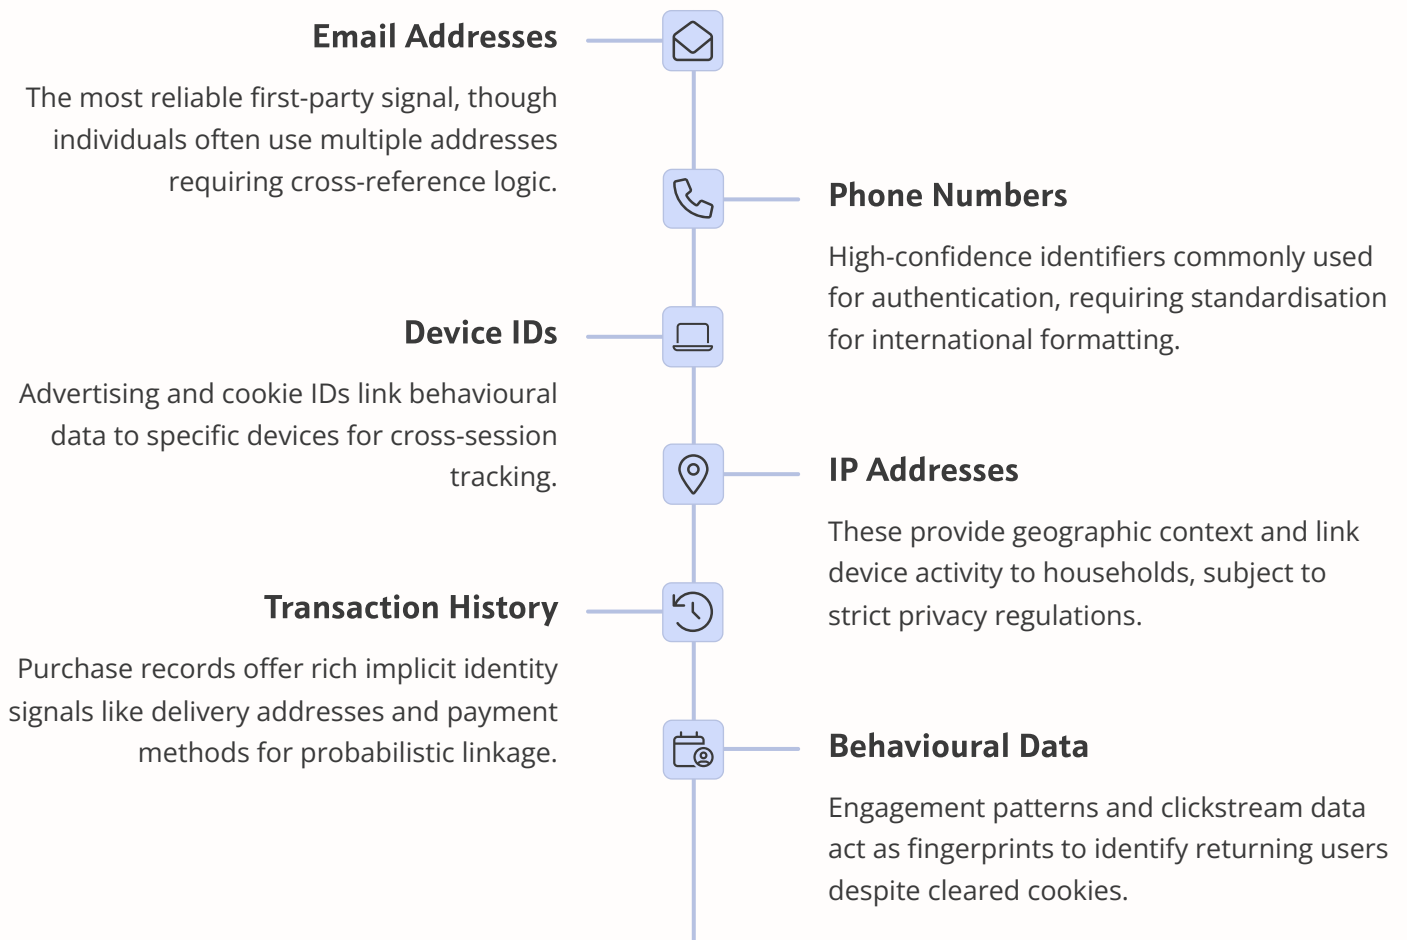
Building an Identity Graph

Creating a robust identity graph requires systematic data modelling, careful infrastructure design, and rigorous governance to manage evolving entity relationships.



Data Sources and Signals

Identity resolution systems rely on the richness and reliability of various identity signals to construct accurate, unified profiles.



Signals vary in reliability, persistence, availability, and privacy sensitivity — a mature strategy maps each signal against these dimensions to govern collection and usage.

- ❑ **Signal Freshness Matters:** An email address collected three years ago may no longer be the primary contact for that customer. Signal recency should be factored into matching confidence scores, with older signals given progressively lower weight.

Privacy, Compliance, and Ethics

Identity resolution demands a balance between powerful data capabilities and a commitment to privacy, transparency, and ethical responsibility.



1

Privacy by Design

Embed privacy into the architecture through encryption, differential privacy, and strict access controls.

2

Consent Architecture

Build a centralised platform to capture, store, and propagate consumer preferences in real time.

3

Audit and Accountability

Maintain logs of linkage decisions and confidence scores to ensure regulatory compliance and debug matching.

Regulatory Requirements

Identity programmes must operate within a complex global regulatory landscape.

- **GDPR (EU):** Requires explicit consent, data minimisation, and right to erasure.
- **CCPA/CPRA (California):** Grants consumers rights to know, delete, and opt out of data sales.
- **HIPAA (US):** Imposes strict controls on linking health-related identity data.

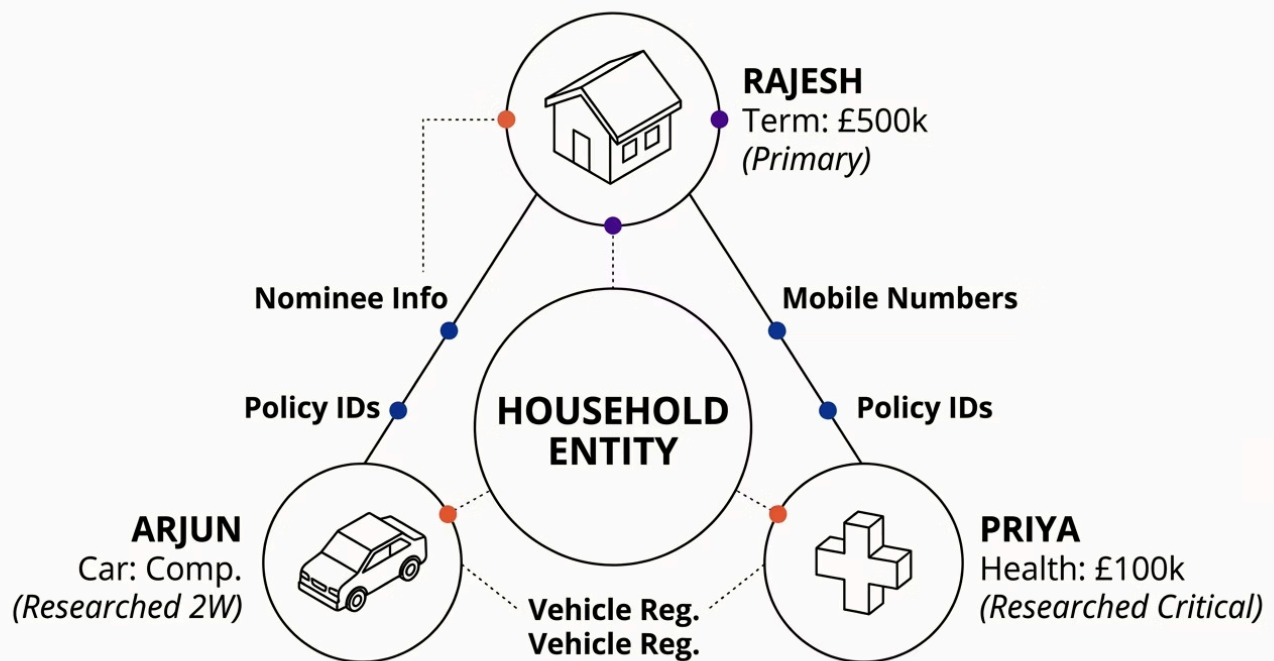
Ethical Principles

Leading organisations apply a higher standard of ethical scrutiny to their programmes.

- **Purpose limitation:** Only process data for clearly communicated and accepted purposes.
- **Data minimisation:** Retain only the signals genuinely necessary for your intended purpose.
- **Transparency:** Ensure individuals understand how their identity data is being used.

Entity-Level Identity Resolution in Action: True 360° Family View

Consider a real insurance scenario: Rajesh Kumar holds term insurance (£500k, 5 years active), his spouse Priya holds a family health floater policy, and their son Arjun holds car insurance. Across traditional systems, these exist as three entirely separate, unlinked profiles. Lemnisk's AI-driven entity resolution — powered by graph neural networks and probabilistic matching — changes everything.



Technical Activation Flow

When Rajesh logs in to check his term policy, the CDP's real-time entity resolution **instantly** identifies all linked family profiles — aggregating household intelligence within milliseconds to inform the next best action.

Unified Strategy Unlocked

- **Coverage gap identified:** Arjun lacks health insurance
- **Income protection gap:** Priya has no income protection policy
- **Pending conversion:** Arjun's two-wheeler quote not yet purchased
- **Offer triggered:** Family Protection Bundle at **25% discount**

📌 This is entity resolution delivering real commercial value: a single login event surfaces a household-level cross-sell opportunity worth thousands in annual premium — made possible only because all three profiles are stitched into one unified family view in real time.

Use Cases Across Industries

Entity-level identity resolution acts as foundational data infrastructure that unlocks value across virtually every sector where customer relationships are central to business performance.



Financial Services

This capability detects suspicious cross-silo fraud patterns while providing a complete view of a customer's institutional relationship. It empowers wealth managers to deliver more relevant advice and effective retention strategies.



Retail & E-Commerce

Identity resolution unifies purchase history and browsing behaviour across in-store and digital channels. This enables truly personalized product recommendations and effective omnichannel campaign targeting.



Healthcare

Unifying patient records across disparate providers reduces duplicate testing and improves medication safety. Accurate linkage is essential for coordinated care and population health management under strict privacy standards.



Travel & Hospitality

Recognising travellers across a complex ecosystem of booking platforms and apps enables seamless, premium guest experiences. It also supports more accurate demand forecasting and targeted ancillary revenue optimisation.

Implementation Roadmap

Implementing a production-grade identity resolution capability requires disciplined sequence and careful programme management.



1. Define Business Objectives



Begin with clear outcomes rather than technology. Identify the specific business problems you are solving to guide every architectural decision.

2. Audit Existing Data



Inventory every system holding customer identifiers and assess overall data quality. This surfaces quick wins and structural dependencies for your plan.

3. Choose Technology Stack



Evaluate build versus buy options based on your technical capabilities and data requirements. Design around open standards to avoid vendor lock-in.

4. Build Matching Models



Develop and validate models using high-quality labelled training data. Establish frameworks that measure both technical performance and business impact.

5. Deploy Identity Graph



Populate your infrastructure with an initial batch of resolved identities. Prioritise high-value segments and validate outputs before expanding coverage.

6. Integrate with Downstream Systems



Connect the graph to CRM, marketing, and analytics platforms. Use robust API contracts and documentation to ensure responsible data usage.


Conclusion

Entity-Level Identity Resolution is now a foundational requirement for delivering coherent customer experiences in our data-complex world. Organizations that prioritize this capability gain a durable competitive advantage over those operating with siloed data.



Start With Business Outcomes

Always anchor your identity resolution program in specific, measurable business objectives.




Invest in Data Quality

Upstream data quality investments are the highest-leverage interventions for any identity resolution program.



Build Privacy In From Day One

Privacy by design is a practical necessity that is far more costly to retrofit later.



Measure Continuously

Use robust measurement frameworks at all levels to drive iterative improvement of the system.

"A well-executed identity strategy enables businesses to unlock the full potential of their data and build lasting customer relationships that create sustainable competitive advantage."

Ready to Build Your 360° Customer View?

You now have the knowledge and strategic perspective to accelerate your identity resolution journey.

Discover Leminisk's Real-Time Entity-Level Identity Resolution and transform fragmented signals into a single, actionable customer truth. Powering smarter decisions, better experiences, and measurable growth.

[Contact Us Today](#)

