

The Economics of Reachability

Why every organisation is paying for a problem it has never measured

A RELAY whitepaper · June 2026
Published for public discussion

Opening Statement

Every organisation on the planet is trying to reach people it can no longer find.

This is not a new problem, but it is a growing one, and it has never been properly named or measured. Organisations have long known that contact data goes stale, that messages sometimes fail to arrive, that customers slip into arrears because a bill did not land, that candidates go unplaced because a recruiter could not get through, that rail passengers are not informed of a cancellation until they are already on the platform. What they have not done is tot up what this costs them – across both the effort required to maintain contact data and the losses that occur when that data fails.

This paper proposes a name, a definition and a framework for that cost: the Cost of Reachability Failure (CRF). It argues that reachability failure is not an operational nuisance to be managed; it is a structural economic category that every organisation already pays, in every sector that depends on reaching people, yet that almost none of them isolates, measures or manages. And it argues that the reason the market has not solved this problem, despite decades of investment in CRM systems, data-quality tools and multichannel campaign engines, is that those tools are all built on the wrong foundation.

Table of Contents

- The Problem: Communication Is Addressed to the Wrong Thing.....1
- The Cost of Reachability Failure (CRF).....2
 - Category 1 – the cost of maintaining reachability.....2
 - Category 2 – the cost of reachability failure itself.....2
- What It Costs: Evidence Across Sectors.....3
 - Recruitment and staffing.....3
 - Energy and utilities.....3
 - Rail and public transport.....3
 - Healthcare.....4
 - Logistics.....4
- The Other Side: Revenue That Reachability Creates.....5
- Why This Problem Cannot Be Fixed – Only Solved.....6
- The Identity Layer.....7
- Conclusion.....8
- About RELAY.....9
- Key Sources.....10

The Problem: Communication Is Addressed to the Wrong Thing

Digital communication has been built almost entirely around endpoints. To reach a person, an organisation stores a pointer to where they were last seen: a mobile number, an email address, a postal address, a messaging handle. That pointer is treated as if it were the person. It is not.

Endpoints are inherently perishable. People change jobs and lose work email addresses. They switch mobile networks and abandon numbers. They move home. They migrate between messaging platforms. Business contact data decays at roughly 2% per month – on the order of 22.5% per year – according to industry analyses repeated across 2024 and 2026 sources including Prospeo, Apollo and ZeroBounce. At any given moment, around 2–5% of records in a large organisation are already outdated.

The result is that every organisation running a large contact database is fighting a war against entropy that it cannot win. It can run faster – with better validation tools, AI-driven enrichment, periodic cleansing – but none of these interventions arrests the decay. They slow it. They make the maintenance treadmill a little smoother. They do not get the organisation off it. And because the decay is permanent and continuous, so is the maintenance cost.

When maintenance fails – when the stored endpoint is wrong, abandoned or simply not attended to – the organisation cannot reach the person. This is reachability failure, and it has a cost. That cost is the subject of this paper.

The Cost of Reachability Failure (CRF)

We propose the Cost of Reachability Failure as a distinct, measurable key performance indicator. It is composed of two cost categories that organisations already pay but never total together.

Category 1 - the cost of maintaining reachability

This is the continuous operational expenditure required simply to keep contact data usable: the collection and re-collection of contact details; validation and verification services; cleansing, deduplication and reconciliation across CRM systems; data-decay management and re-enrichment cycles; and the staff hours consumed in all of the above. Validity's State of CRM Data Management in 2025, based on 602 CRM users and administrators across the US, UK and Australia, found that frontline workers spend on average 13 hours per week hunting for basic contact information, and that 76% of organisations say less than half of their CRM data is accurate and complete. This is not a marginal inefficiency. It is a permanent structural tax on every organisation that stores contact endpoints.

Category 2 - the cost of reachability failure itself

This is the downstream loss that occurs when a person cannot be reached at all: the vacancy left open because a candidate could not be contacted; the customer who slips into arrears because a bill or reminder never arrived; the undelivered message that becomes a second and third contact attempt, each carrying its own unit cost; the missed appointment that wastes clinical capacity; the passenger who was not informed of a cancellation and generates a compensation claim; the package that could not be delivered and triggers a re-attempt, a return, and potentially a lost customer. These are the bookable losses that organisations currently mis-attribute to "operations" or "bad luck" rather than to reachability.

Why CRF is unmeasured. The maintenance costs sit in IT and data-governance budgets; the failure costs are dispersed across recruitment, collections, customer service and compliance. No single owner sees the total, and so no one manages it. CRF is the metric that makes the total visible.

CRF per person per year = (annual maintenance cost allocated to that record) + (probability of reachability failure × cost per failure event × failure events per year)

What It Costs: Evidence Across Sectors

The evidence for CRF is available across every communication-intensive sector. The figures below are drawn from recent, publicly available sources and are presented with precise descriptions of what each figure actually represents. Where a number is a broad aggregate rather than a reachability-specific cost, that is stated explicitly.

Recruitment and staffing

Staffing is a reachability business in its purest form. The product being sold – a successfully placed candidate – can only be delivered if the right person can be reached, confirmed and deployed within a narrow commercial window. SHRM's 2025 Talent Acquisition Benchmarking Report, based on 2,371 member responses, puts the average cost per hire for non-executive roles at USD 5,475, with an average time-to-fill of 44 days. Every day a vacancy stays open carries a productivity cost commonly estimated at around USD 500 for a standard role. And every candidate who cannot be reached for redeployment – a second or third assignment without fresh sourcing cost – forces the agency to pay the full cost-per-hire again, for someone it has already sourced, vetted and placed. Industry analysis suggests top-performing light-industrial agencies achieve redeployment rates of 40–60%, while agencies below 25% are in effect throwing away acquisition spend. The difference between those bands is largely a function of whether the agency can reach the worker at the right moment – which, after a gap of even a few weeks, is a reachability problem.

Energy and utilities

UK domestic consumer energy debt reached £4.43 billion by June 2025, an increase of 71% since 2023, with 2.44 million consumers in debt and nearly three-quarters of that total held by customers with no repayment plan in place (Ofgem, Debt Strategy Update, November 2025). The link to reachability is direct: customers who cannot be reached for payment-arrangement conversations are the ones who drift from manageable arrears into write-offs. Ofgem requires suppliers to make at least ten documented contact attempts before installing a prepayment meter; that obligation is impossible to discharge against a stale address. On the maintenance side, the UK's regulated price cap allows roughly £97 per year in core operating costs and £26 per year in debt-related costs per direct-debit customer (House of Commons Library, 2026, citing Ofgem) – a significant portion of which is consumed by communication and contact-management activity. Per Ofgem's own analysis, a typical consumer pays around £52 per year towards the cost of managing and writing off energy debt – a cost that better reachability would materially reduce.

Rail and public transport

The rail sector illustrates the reachability problem with unusual clarity, not least because it is where the Cost of Reachability Failure thesis first crystallised. A railway operator maintains a continuous relationship with only a small fraction of its passengers – season-ticket and subscription holders – while the majority travel infrequently. In the months between journeys, a passenger's phone number, email address and channel preferences may all have changed. When disruption strikes, the operator has minutes – not days – to reach the right person. Reaching them on a dead channel five minutes too late is operationally identical to not reaching them at all.

The regulatory and financial consequences are substantial. Under Regulation (EU) 2021/782 on rail passengers' rights, operators must compensate passengers for delays of 60 minutes or more (25% of the ticket price at 60–119 minutes; 50% at 120 minutes or more) and provide real-time information during disruption. Deutsche Bahn, whose long-distance punctuality fell to 62.5% in 2024 – the worst in at least 21 years – paid 197 million euros in passenger compensation in 2024, against 6.9 million claims (DB Facts & Figures 2024; dpa/iamexpat, March 2025). Every compensation payment that could have been avoided by reaching the passenger in time, and every claim that could have been pre-empted by proactive communication, is a direct expression of CRF.

Healthcare

NHS England estimates that around 7.2 million GP appointments are missed each year, at a cost of approximately £216 million – equivalent to the salaries of 2,325 full-time GPs (NHS England, January 2019). Analysis of 2024 Freedom of Information data by Esendex estimated 11.8 million missed hospital appointments across NHS trusts at a total cost of around £1.9 billion. Among the most commonly cited reasons for non-attendance: the patient was unaware of the appointment, forgot it, or received unclear or incorrect information. These are reachability failures. Each missed appointment is a waste of clinical resource and a delay of care for someone else – and it is directly preventable through reliable contact with the patient at the right time on the right channel.

Logistics

Per Loqate's Fixing Failed Deliveries study (2021), around 8% of domestic first-time deliveries fail, costing retailers an average of USD 17.20 per order, with 71% of businesses citing inaccurate delivery addresses as a primary cause. More recent analysis reports first-attempt failure rates of 5–15% in normal conditions, rising to as high as 20% at peak. Seventy per cent of consumers who experience a failed delivery say they are unlikely to return to that retailer (SmartRoutes, 2025). The cost of a failed delivery is therefore not just the re-attempt; it is the customer relationship that was lost because a message – a simple "you are not home, here is how to reschedule" – did not reach the right person in time.

The Other Side: Revenue That Reachability Creates

The CRF framework focuses on cost, but its mirror image is equally important. Reliable reachability does not merely reduce losses – it creates value that would otherwise not exist.

In staffing, the redeployed contractor is pure margin: no sourcing cost, no job-board fee, no agency process – just the assignment fee on a worker the agency has already paid to acquire. In energy, the customer reached early for a payment arrangement becomes a recoverable balance rather than a write-off; the vulnerable customer reached for social-tariff enrolment becomes a supported, retained relationship rather than a complaint and a churn event. In logistics, the customer who receives a proactive notification after a missed delivery is 14% more likely to place a second order within 90 days than one who receives no communication at all (DeliverLoop analysis); the failed delivery becomes a non-event rather than a lost customer. In banking, the proactive fraud alert – which only works if the bank can reach the customer in the seconds that matter – prevented £1.45 billion of unauthorised fraud in 2024, equivalent to 67p in every £1 attempted (UK Finance, 2025).

These are not soft benefits. They are measurable revenue events – fees earned faster, debts recovered, customers retained, fraud prevented – each of which depends on a single precondition: that the right person was reached at the right moment.

Why This Problem Cannot Be Fixed - Only Solved

If reachability failure is so costly, and the costs are so well evidenced, why has the market not solved it? Organisations already spend heavily on CRM systems, data-enrichment services, phone-number verification APIs and email-validation tools. The honest answer is that all of these tools are addressing the wrong thing. They are endpoint-centric: they try to keep up with the decay of stored contact details rather than removing the need to store them. And in a race against continuous, exogenous entropy, endpoint-centric tools can at best slow the rate of loss. They cannot stop it.

This is not a criticism of those tools. It is a structural observation about the paradigm they are built on. A better CRM stores endpoints more elegantly; AI-driven enrichment guesses at fresher endpoints from more data; a verification API confirms whether an endpoint is currently live. All of them still operate on the endpoint as the unit of address – and so all of them still inherit the endpoint's fatal property: it is not the person, and it decays.

The organisation's contact database is not a stable record of who its customers are. It is a perishable snapshot of where they were when last seen.

The only structural solution is to stop addressing endpoints and start addressing people. This is the same move the internet has made twice before. The Domain Name System solved internet routing by decoupling the stable name from the perishable IP address: DNS introduced a layer of abstraction above the address, so that the address beneath could change freely while the name endured. Email performed the same move for messaging, giving each person a portable identity that survived changes of internet service provider. But email's abstraction was only partial: the email address itself has become perishable, tied to an employer, institution or platform that the person will eventually leave.

The relay address is the next abstraction. It sits above any particular number, email address or messaging platform, and it is maintained by the person – not the organisation. When the person changes their number, their email, or their preferred app, they update the routing behind their identity once, on their own side. Every organisation that addresses their identity continues to reach them, with no cleaning, no enrichment and no re-verification, because none of those organisations was ever storing the endpoint in the first place.

The Identity Layer

RELAY is a communication infrastructure layer, not an application or a messaging service. Every person, organisation and AI agent receives a persistent relay address – for example name.relay – as a single, durable communication endpoint. The sender needs to know only the identity; RELAY routes the message to wherever and however the recipient currently chooses to receive it, across whatever channels they prefer, without the sender needing to know or maintain any of that.

For an organisation, this collapses the CRF maintenance category: there are no endpoints to store, validate or cleanse, because the routing is maintained by the recipient. And it compresses the failure category: messages addressed to the identity reach a live, currently active channel rather than a stale stored number or an abandoned inbox.

The analogy that best captures the relationship is the domain name and the server. A domain name points to a server, and the server's IP address can change without breaking the name. A relay address points to a person, and the person's channels can change without breaking the address. The sender, the organisation, the regulator, the rail operator trying to reach a delayed passenger in real time – they all address the identity, and the identity never decays.

Conclusion

Reachability is infrastructure. Like all infrastructure, its cost is invisible until it is isolated and measured. The Cost of Reachability Failure framework proposes a way to make that cost visible – across sectors, across organisations, and at the level of a single customer record. The evidence presented in this paper suggests that CRF is one of the largest unmeasured cost categories in modern enterprise operations, running through staffing, energy, banking, rail, healthcare, logistics and every other sector that depends on regular contact with people.

The organisations that measure it first will be the ones that recognise, before their competitors, that the cheapest message is the one that does not have to be sent twice – and that the most valuable message is the one that reaches exactly the right person at exactly the right moment. CRF is the metric that makes the difference between those two outcomes visible. The identity layer is the infrastructure that closes it.

RELAY is the layer that lets them.

About RELAY

RELAY is a communication infrastructure layer in development, built around a persistent identity model: every person, organisation and AI agent receives a relay address that functions as a durable, channel-agnostic communication endpoint, with the recipient controlling routing across channels. RELAY is identity-driven, user-controlled and channel-agnostic. It is not an application, an inbox or a messaging service.

RELAY is currently in the feasibility and early-adoption phase, operating through SynergyPeak & Company (KvK 92665241) under founder and CEO Jasper van Berkel. A dedicated entity will be incorporated upon first funding. The full research report on which this whitepaper is based, The Economics of Reachability: A RELAY Research Report on the Cost of Reachability Failure, is available for investor and enterprise discussions.

For enquiries: jasper@synergypeakcompany.com

Key Sources

The quantitative claims in this whitepaper are drawn from the following sources. Full references, including precise descriptions of what each figure represents and explicit labelling of RELAY's own working estimates, are set out in the accompanying full research report.

- Aqore, Temp worker management software: what it is and what to look for in 2026 (redeployment rate benchmarks).
- ContactBabel, UK Contact Centre Decision-Makers' Guide 2024; US Contact Center Decision-Makers' Guide 2025.
- DB Facts & Figures 2024; iamexpat.de / dpa (March 2025) – €197m passenger compensation, 6.9m claims.
- DeliverLoop, The Real Cost of Missed Delivery Attempts (repeat-purchase impact).
- Esendex, Counting the cost of missed NHS appointments (2024–2025).
- House of Commons Library, What costs make up an electricity bill? (2026, citing Ofgem).
- Loqate, Fixing Failed Deliveries (Censuswide, 2021).
- NHS England, missed appointment cost data (January 2019).
- Ofgem, Debt strategy update: supporting the reduction of energy debt (November 2025) and response to ESNZ Committee.
- Office of Rail and Road, Rail Fares Index 2024.
- Prospeo, B2B Contact Data Decay in 2026; Apollo; ZeroBounce (2025) – data-decay rates.
- Regulation (EU) 2021/782 on rail passengers' rights and obligations (applicable June 2023).
- SHRM, 2025 Benchmarking (Talent Access) Report.
- SmartRoutes, Delivery Success Rates: Key Stats for Retail and eCommerce (2025).
- UK Finance, Annual Fraud Report 2025.
- Validity, The State of CRM Data Management in 2025.