THE ECONOMIC AND SOCIETAL COST OF INFRASTRUCTURE - LAYER INACTION

Author: Robert Lee
Date: June 2026

Executive Summary

This paper is the third in a four-part series examining the future of telecommunications resilience. The first paper established the technical reality of cross-domain, AI-accelerated threats. The second diagnosed the structural regulatory failure that has left those threats largely invisible to oversight. This paper does something different: it puts numbers against the problem.

The cost of infrastructure-layer inaction is not abstract. It is measurable across four dimensions, carriers, regulators, enterprise customers, and retail consumers, and it is compounding every year that structural frameworks are not updated. Drawing on documented breach events, industry cost data, regulatory enforcement records, and direct operational experience of what infrastructure-level incidents do to businesses in practice, this paper presents the most comprehensive accounting of infrastructure-layer harm costs yet assembled.

The headline findings are significant. Global telecom fraud losses exceeded US$38.95 billion in 2023 and are projected to grow by 15 to 20 percent annually without structural intervention. The average cost of a data breach in 2025 reached US$10.22 million, and for major carrier incidents, the true total cost when pipeline loss, customer churn, operational lockdown, and long-term reputational damage are included is a multiple of that figure. GDPR enforcement against telecoms is tightening, with notification windows reduced to 48 hours from 2025, while the tools to meet those obligations remain largely manual and retrospective.

But this paper also presents a more optimistic argument, one that distinguishes it from a standard cost assessment. The same structural change that would reduce these costs also generates its own funding mechanism. Automated, evidence-grade infrastructure-layer oversight creates a short-term enforcement revenue surge for regulators that funds the transition, before evolving into an industry-wide self-regulation model that removes the need for costly reactive enforcement altogether. The cost of inaction is high. The cost of action is lower, and largely self-financing.

Paper Four in this series will set out the operational model for implementing this transition.

$38.95B	$10.22M	110M
Global Telecom Fraud Losses	**Average Data Breach Cost**	**AT&T Customers Compromised**
2023 - growing 15-20% annually	2025 IBM report - US market highest globally	Single 2024 breach - call & text records stolen

48hrs	6.4M	40%
GDPR Notification Window	**Bouygues Customer Records**	**Breach Notification Recipients**
Telecoms from 2025 - down from 72 hours	August 2025 - IBANs and financial data exposed	Subsequently targeted in downstream phishing

NUMBERS AT A GLANCE

The Cost of Invisibility

The first two papers in this series established a structural paradox: infrastructure-layer harm is among the most consequential threat category facing the UK digital economy, yet it is the least monitored, least reported, and least enforced. This paper examines what that invisibility costs in practice.

Standard breach cost analyses, including the widely cited IBM Cost of a Data Breach Report, capture direct costs reasonably well: forensics, legal fees, regulatory notification, immediate remediation. What they consistently understate are the indirect and long-term costs: the sales pipeline that freezes during a six-month operational recovery, the customers who chose a competitor during that period, the insurance premiums that permanently reset upward, the staff productivity diverted from revenue generation to crisis management, and the reputational damage that compounds quietly in B2B procurement decisions for years after the incident is technically resolved.

For telecom carriers specifically, whose business model depends on long-term contracted relationships with enterprise customers who require absolute confidence in network integrity, these indirect costs are not secondary. They are often the dominant cost. A carrier managing a ransomware recovery for six months is not merely spending on incident response. It is haemorrhaging the contracted recurring revenue pipeline that it will never recover, because enterprise procurement cycles are long and institutional memory is longer.

This paper attempts a more complete accounting. The costs documented here span four stakeholder groups, carriers, regulators, enterprise customers, and retail consumers, and include both the direct costs that appear in breach disclosures and the indirect costs that rarely do. The aggregate picture is significantly larger than headline figures suggest.

The Cost to Carriers

Carriers bear the most immediate and direct costs of infrastructure-layer harm, but also the most underreported ones. The gap between the costs that appear in regulatory disclosures and the true commercial impact of a major incident is substantial.

Direct Breach Costs

The direct costs of a major carrier breach are well-documented in aggregate but rarely disclosed in full by individual operators. Forensic investigation, legal counsel, regulatory notification, technical remediation, third-party security assessment, and customer notification combine to represent a significant immediate outlay. IBM's 2025 Cost of a Data Breach Report puts the average global breach cost at US$4.88 million, rising to US$10.22 million in the United States, the highest of any country globally. For a carrier operating across multiple jurisdictions with GDPR obligations and telecommunications-specific regulatory requirements, the compliance and notification costs alone can exceed this average.

The Colt Technology Services incident of August 2025 illustrates the scale at the carrier level. The WarLock ransomware group claimed exfiltration of one million documents including financial records, employee salary data, executive communications, and customer contact details. The Colt Online customer portal, Voice API platform, and a range of internal support systems were taken offline for multiple days, with full restoration of automated monitoring capability taking significantly longer. The regulatory notification obligations alone, across multiple jurisdictions in Europe, Asia, and North America, represent a substantial immediate compliance cost.

The Operational Lockdown Cost, What Breach Disclosures Don't Capture

What standard breach cost analyses do not capture is what

operational near-paralysis does to a carrier business over an extended recovery period. A major carrier incident does not end when the forensic team declares the immediate threat contained. The downstream operational impact can persist for months, and its commercial consequences are severe.

The Operational Reality of a Major Carrier Incident

When a carrier of Colt's scale and profile suffers a major cyber incident, the commercial impact extends far beyond the immediate remediation costs. New enterprise sales effectively stop, no procurement team responsible for a multi-year managed network contract will sign during an active incident recovery. The sales pipeline freezes at precisely the moment the business is incurring its highest costs. Existing customers invoke SLA review clauses. Legal and compliance teams divert capacity from commercial activity to incident response. The automated systems that underpin efficient network operations revert to manual processes, increasing operational cost while reducing service quality. The organisation is, in effect, fighting a war on multiple fronts simultaneously, technical, legal, commercial, and reputational, while its revenue generation capability is severely constrained.

The commercial cost of six months of constrained operations for a carrier of this scale, factoring in frozen pipeline, customer churn risk, SLA penalties, staff productivity diversion, and the loss of competitive positioning during the recovery period, is likely to represent a multiple of the direct incident costs. This figure almost never appears in public breach disclosures, but it is the cost that matters most to carrier boards and investor committees.

The Hidden Pipeline Cost

Perhaps the most under-appreciated cost of a major carrier incident is the pipeline it never generates. Enterprise customers evaluating carrier options during a competitor's incident recovery do not pause their procurement process, they redirect it. The contracts signed with a competitor during that period are not merely deferred: they are typically lost for the duration of the contract term, which in enterprise carrier agreements is commonly three to five years. For a carrier managing a significant volume of enterprise procurement activity, the pipeline value lost during a six-month recovery period can substantially exceed the direct costs of the incident itself.

This pipeline loss does not appear in breach cost disclosures. It does not appear in regulatory filings. It is invisible in the public record. But it is real, it is large, and it compounds with every month that recovery takes longer than customers will tolerate.

Insurance Premium Reset

A carrier that has suffered a major cyber incident will face a significant upward reset in its cyber insurance premiums at renewal. This is not a one-time cost, it is a permanent increase in the carrier's cost base that persists for years after the incident. For larger carriers with significant cyber insurance exposure, the premium increase following a major breach can represent millions of pounds of additional annual cost, indefinitely. This cost, like pipeline loss, rarely appears in breach cost analyses but is a real and lasting financial consequence of infrastructure-layer inaction.

The AT&T Scale Comparison

For international context, the AT&T breach of 2024 illustrates infrastructure-layer costs at hyper-scale. A cyberattack on AT&T's cloud data environment compromised data on almost all of the carrier's wireless customers, approximately 110 million people, with attackers accessing call and text message records spanning six months. AT&T reached a US$13 million settlement with the FCC for a separate 2023 breach affecting 8.9 million customers. The combined regulatory, legal, and remediation costs across AT&T's multiple 2024 breach events represent one of the largest carrier-level incident cost profiles on record, yet even these figures are likely to understate the true commercial impact when indirect costs are included.

Cost Category	Direct Cost	Hidden Cost	Compounding Factor
Direct remediation	Forensics, legal, notification	Diversion of senior leadership time	Repeats with every subsequent incident
Operational disruption	Manual process costs, SLA penalties	Staff productivity loss across organisation	Compounds with recovery duration
Pipeline freeze	Deals on hold during recovery	Contracts signed by competitors	Lost for full contract term (3–5 years)
Customer churn	Immediate contract exits	Reduced renewal rates post-incident	LTV loss per churned customer
Insurance premium reset	Higher renewal premium	Permanent cost base increase	Grows with each subsequent incident
Reputational damage	PR and communications spend	Procurement score reduction	B2B trust deficit persists for years

The Cost to Regulators

The cost of infrastructure-layer inaction to regulators is different in character from the costs borne by carriers or customers, but it is no less significant. It manifests in three ways: the direct resource cost of manual, retrospective enforcement; the enforcement revenue that cannot be collected because the evidence framework does not exist; and the institutional credibility cost of documented blind spots.

The Resource Cost of Reactive Enforcement
Infrastructure-layer enforcement under the current framework is manual, domain-specific, and retrospective. Building a prosecutable case against a carrier for network-level misuse requires Ofcom investigators to manually assemble evidence across disconnected data sources, coordinate with multiple operators, and navigate the fragmented multi-body mandate described in Paper Two. A single enforcement action of meaningful complexity can consume months of regulatory resource. At a time when Ofcom's attention is, as Paper Two documented, already stretched across online safety, broadcast regulation, consumer protection, and spectrum management, this is a significant and avoidable cost.

Automated, evidence-grade infrastructure-layer telemetry changes this equation fundamentally. When the evidence file is assembled automatically, timestamped, immutable, and built to the evidentiary standard required for legal proceedings, the regulatory resource required per enforcement action falls dramatically. More cases can be pursued. Enforcement velocity increases. And the deterrent effect of visible, consistent enforcement begins to change carrier behaviour in ways that reduce the volume of future cases.

The Enforcement Revenue Gap, Fines Left Uncollected
This is perhaps the most provocative argument in this paper, and the most financially significant for regulators. Ofcom holds enforcement powers that carry substantial financial penalties.

The Enforcement Revenue Argument

An automated intelligence platform that generates evidence-grade case files against non-compliant carriers and ISPs does not merely reduce regulatory costs, it unlocks enforcement revenue that the current framework leaves entirely uncollected. The volume of prosecutable infrastructure-layer violations that would become visible under a unified telemetry model is likely to be substantial. Early enforcement actions would generate significant penalty revenue, revenue that funds the platform, reduces the regulatory resource burden, and creates the deterrent effect that drives subsequent compliance. This is not a cost to the regulator. It is a revenue opportunity that the current blind spot is preventing them from realising.

The GDPR Tightening, Rising Obligations, Static Capability
From 2025, GDPR notification timelines for telecoms were reduced from 72 to 48 hours for significant breach events. This tightening of obligations is occurring at the same time as the volume and complexity of infrastructure-layer incidents is increasing. Regulators are being asked to receive, process, and act on more notifications, faster, with the same manual tools and the same fragmented reporting infrastructure that existed before. The gap between regulatory obligation and regulatory capability is widening, and the cost of that gap, in terms of enforcement failures and reputational exposure, is growing with it.

The Institutional Credibility Cost
There is a less quantifiable but genuinely significant cost to regulators in the persistence of documented blind spots. When a major carrier breach, such as the Colt incident of August 2025, or the Salt Typhoon campaign that compromised nine US telecommunications providers over a two-year period, demonstrates that infrastructure-layer harm was occurring at scale without regulatory detection or intervention, the question of whether the regulator was adequately equipped to fulfil its mandate becomes unavoidable. For Ofcom, which is accountable to Parliament and to the public for the safety and integrity of UK communications infrastructure, this is not merely a reputational concern. It is a governance one.

The Cost to Enterprise Customers

Enterprise customers — the businesses that depend on carrier infrastructure for their communications, data connectivity, and operational systems — bear significant costs from infrastructure-layer harm, most of which are invisible in carrier breach disclosures. When a carrier is compromised, its enterprise customers do not merely experience service disruption. They face direct security exposure, operational disruption, and in some cases regulatory liability of their own.

The Salt Typhoon Benchmark, State-Level Carrier Exploitation
The Salt Typhoon campaign, identified in late 2024, represents the most significant documented example of infrastructure-layer carrier exploitation affecting enterprise and government customers. The Chinese state-sponsored threat actor infiltrated at least nine major US telecommunications providers, including AT&T, Verizon, T-Mobile, and Lumen Technologies, over a two-year period, accessing call records, geolocation data, and in some cases the private communications of individuals involved in government and political activity.

The campaign exploited well-documented vulnerabilities that had been catalogued in CISA's Known Exploited Vulnerabilities register but not patched.

For enterprise customers of the affected carriers, the implications were profound. Business communications that transited the compromised carrier networks were potentially accessible to a state-level adversary for an extended period. The organisations whose data was exposed had no visibility into this exposure, it was not disclosed in real time, could not be detected through their own security tooling, and was only identified through national intelligence investigation. This is infrastructure-layer harm at its most consequential: entirely invisible to the enterprise customer, operating inside the trusted carrier relationship, and exploitable over an extended period precisely because no cross-domain detection capability existed.

AI Voice Fraud and Executive Impersonation
At a more granular but rapidly growing scale, AI-generated voice fraud targeting enterprise customers represents one of the fastest-growing categories of carrier-enabled harm. Deepfake audio technology has reached the point where executive voice impersonation is commercially viable at industrial scale. Fraudsters use AI-generated voice calls, routed through carrier infrastructure, exploiting CLI spoofing, and designed to defeat biometric authentication, to authorise fraudulent financial transfers, extract credentials, or gain access to sensitive systems. The carrier is the infrastructure through which these attacks are delivered. Without cross-domain intelligence at the carrier layer, individual enterprise security tools cannot detect them.

Cascading SLA and Operational Impact
When a carrier suffers a major operational incident, whether through external attack, infrastructure failure, or the kind of extended recovery period that follows a significant breach, enterprise customers bear direct operational costs. SLA failures trigger credit claims and contract review processes. Critical business systems that depend on carrier connectivity may be degraded or unavailable. Business continuity plans are invoked, diverting management attention and incurring additional cost. For enterprise customers operating real-time processing environments, financial trading, logistics management, manufacturing control systems, even brief carrier degradation events can have disproportionate operational and financial consequences.

The Enterprise Cost of AI Voice Fraud

A 2024 case involving a multinational financial services firm illustrates the enterprise cost dimension. Fraudsters used AI-generated audio to impersonate the firm's CFO in a voice call to the finance team, authorising a transfer of approximately US$25 million to overseas accounts. The call transited carrier infrastructure using a spoofed CLI presenting as the CFO's known number. No carrier-layer detection flagged the anomaly. The firm's own security tooling had no visibility into the call routing chain. By the time the fraud was identified, the funds had been moved through multiple jurisdictions. Recovery was partial. The reputational and regulatory consequences for the firm, which had an obligation to maintain adequate fraud controls, were significant.

The Procurement Consequence
An increasingly important and under-appreciated consequence of high-profile carrier incidents is their effect on enterprise procurement behaviour. Chief Information Security Officers and procurement teams at major enterprise customers are increasingly factoring carrier security posture into vendor selection decisions. A carrier with a documented breach history, or one that cannot demonstrate cross-domain visibility and evidence-grade incident response capability, faces a growing procurement disadvantage. This is a commercial consequence for carriers, but for enterprise customers it represents a market signal that the baseline of carrier security expectation is rising, and that carriers unable to meet it will lose business to those that can.

The Cost to Retail Consumers

Retail consumers are the ultimate downstream recipients of infrastructure-layer harm. They are also the least able to protect themselves from it, because by definition infrastructure-layer threats operate below the visibility threshold of any individual consumer security tool.

Financial Fraud - The Consumer Cost
The UK fraud landscape provides the starkest quantification of consumer-level harm. Annual fraud losses in the UK exceed £2.3 billion, with a significant and growing proportion enabled at the telecommunications infrastructure layer, through CLI spoofing, smishing at scale, AI-generated voice fraud, and the grey-route traffic that enables fraudulent calls to reach consumers at low cost. Ofcom's own data documents 45 million scam calls in the UK in a single reporting period. The majority of these calls exploit carrier infrastructure, using spoofed caller identity, unmonitored interconnect routes, and machine-generated traffic volumes, in ways that are detectable at the network layer but invisible to the consumer receiving the call.

The financial impact on individual consumers ranges from modest, a scam call that wastes time but results in no financial loss, to catastrophic, with authorised push payment fraud and

investment scams regularly resulting in individual losses of tens of thousands of pounds. The aggregate consumer harm is substantial: UK Finance data consistently shows that fraud is the most common crime experienced by adults in England and Wales, with telecoms-enabled fraud representing a significant and growing component.

Compromised Home Networks - The IPTV Device Threat

Paper One documented in detail the threat posed by pre-infected illegal IPTV devices, streaming boxes and set-top devices that arrive with malware pre-installed, connecting consumer home networks to criminal botnet infrastructure. The scale of this threat is significant: 216 billion visits to illegal streaming sites were recorded globally in a single reporting period, with the associated device infrastructure representing one of the largest vectors for consumer network compromise currently active.

For the retail consumer, the consequences of a compromised home device extend well beyond the illegal streaming service they believed they were purchasing. Their home network becomes a node in criminal infrastructure. Their connected devices, smart home systems, personal computers, mobile devices on the same network, are potentially accessible to the same criminal actors. Their broadband data is consumed by botnet activity. And they have no visibility into any of this, because it operates entirely at the network layer, below the threshold of any consumer-facing security tool.

The Downstream Phishing Effect

Research consistently shows that individuals whose data has been exposed in a breach are significantly more likely to be subsequently targeted by phishing attacks, with one analysis finding that approximately 40 percent of breach notification recipients experience targeted follow-up phishing attempts. For retail consumers who receive breach notifications from their carrier, as millions did following the Bouygues breach of 6.4 million customers in 2025, the AT&T breach affecting 110 million customers, and the Odido breach affecting 6.2 million Dutch consumers in early 2026, the breach notification is not the end of the harm. It is frequently the beginning of a second wave.

The Erosion of Trust

Beyond the direct financial and security costs, infrastructure-layer harm imposes a broader societal cost through the erosion of consumer trust in digital communications. When consumers cannot trust that a phone call presenting a known number is genuine, cannot trust that a carrier-authenticated SMS has not been spoofed, and cannot trust that the streaming device they purchased does not embed malware, the foundational trust that underpins digital participation is damaged. This is difficult to quantify but significant in its consequences, for the adoption of digital public services, for the functioning of digital commerce, and for the broader social contract between citizens and the digital infrastructure they depend on.

The Reputational Cost - A Dimension That Multiplies Everything Else

Reputational cost is often treated as a soft metric, difficult to quantify, easy to dismiss in a financial analysis. For telecommunications carriers, this is a serious analytical error. Reputation is not peripheral to a carrier's commercial model. It is central to it.

Why Reputation Matters More for Carriers Than Most Industries

Enterprise and wholesale carrier relationships are built on trust at a structural level. A carrier that underpins banking networks, data centre connectivity, critical national infrastructure, and enterprise communications for thousands of organisations is not selling a commodity service, it is being trusted with the operational continuity of those organisations. When that trust is broken, the reputational damage does not merely affect future sales. It infects existing relationships, triggers contract review processes, and reduces the carrier's negotiating position on every renewal it faces in the years following a major incident.

The Transparency Failure Multiplier

The Colt incident of August 2025 illustrates a pattern that is consistent across major carrier breaches: the reputational damage from the incident itself is often compounded by the communications response to it. Colt initially described its multi-day service disruption as a technical issue with no data loss, before subsequently acknowledging potential file exfiltration after the WarLock ransomware group published data on its leak site. The gap between initial assurances and subsequent disclosures, a gap that security researcher Kevin Beaumont publicly highlighted, generated a second wave of reputational damage that was in some respects more damaging than the original incident. Enterprise customers and industry observers noted not merely that Colt had been breached, but that its initial communications could not be relied upon. For a carrier whose value proposition rests on trustworthiness, this is a serious and durable reputational injury.

The Board and Investor Consequence

Major carrier incidents have consequences that extend well beyond the operational and commercial layer. Board-level governance scrutiny increases following a significant breach, with non-executive directors facing questions about whether adequate oversight of cyber risk was maintained. For publicly listed carriers, share price impact is measurable and often sustained. Investor relations teams face difficult conversations about the adequacy of pre-incident security investment and the likely cost of post-incident remediation. These governance and investor consequences represent a real and quantifiable cost that, again, rarely appears in standard breach cost analyses.

The Compounding Transparency Problem

When a carrier initially minimises a breach and is subsequently shown by independent reporting or threat actor disclosures to have understated its severity, the reputational damage compounds in three ways simultaneously. Customer trust in the carrier is damaged directly. Trust in the carrier's crisis communications, and by extension its normal operational communications, is damaged structurally. And the incident becomes a reference point that procurement teams at enterprise customers will cite in future contract negotiations, security questionnaires, and vendor assessments for years. The initial breach is a recoverable event. The transparency failure that follows it is often not.

The Long-Term Brand Recovery Cost

The cost of recovering from reputational damage at the scale associated with a major carrier breach is substantial. Marketing spend, thought leadership investment, third-party security certification, and the extended period during which the carrier must demonstrate sustained operational integrity before enterprise procurement committees are willing to trust it again, all represent ongoing costs that persist long after the technical incident has been resolved. For carriers that compete in markets where trust is a primary differentiator, this long-term brand recovery investment can represent tens of millions of pounds over a multi-year period.

The Self-Financing Model - How Action Funds Itself

The most important argument in this paper is not the cost of inaction, though that cost is substantial and documented. It is the argument that the structural change required to address infrastructure-layer harm is largely self-financing. Understanding why requires following the flywheel logic through its five stages.

The Flywheel Model

The deployment of automated, evidence-grade infrastructure-layer intelligence does not impose a net cost on the regulatory system. It generates one of the most powerful economic mechanisms available to a regulatory body: the ability to collect penalties that were always legally available, from violations that were always occurring, but that could never previously be evidenced and prosecuted at scale.

The Answer to 'Who Pays?'

Every regulatory reform proposal faces the same objection: who pays for it? The flywheel model provides a precise and defensible answer. The non-compliant carriers pay for it, through the fines they should always have been paying, but could not be made to pay because the evidence framework did not exist. The enforcement revenue generated in the early stages of platform adoption funds the intelligence infrastructure that generates it. The regulator does not need to find new budget. Parliament does not need to allocate new resource. The market funds its own correction.

This is not a theoretical model. It is the logic that underpins the most successful regulatory enforcement frameworks in comparable sectors. Financial services regulators have used automated transaction monitoring to generate enforcement revenue that funds supervision of the very firms being monitored. Environmental regulators have deployed real-time emissions monitoring to unlock penalty income that was previously invisible. The telecommunications sector is the outlier, the sector where the infrastructure layer has remained unmonitored and where the enforcement revenue that should flow from that monitoring has remained uncollected.

The Short-Term Boost and the Long-Term Dividend

It is important to be clear about the temporal structure of this argument. The enforcement revenue boost is a short-term phenomenon, it is highest in the period immediately following platform adoption,

Stage	What Happens	Financial Outcome
Stage 1	Regulator adopts unified intelligence platform. Automated evidence-grade case files generated against non-compliant carriers and ISPs.	Enforcement actions that previously took months now take days. Volume of prosecutable cases increases dramatically. Penalty revenue begins to flow.
Stage 2	Enforcement revenue surge funds regulatory operations and enables resource reallocation from manual investigation to higher-value oversight activity.	Ofcom redirects capacity toward online harm agenda. Platform self-finances through the penalties it enables. Regulatory capability improves while cost base stabilises.
Stage 3	Carriers and ISPs face evidence-grade, timestamped visibility into their own non-compliance. Fine risk becomes existential for smaller operators, reputationally catastrophic for larger ones.	Commercial incentive to self-regulate becomes overwhelming. Operators adopt the platform to avoid enforcement — not out of altruism, but out of financial self-preservation.
Stage 4	Carriers with platform capability gain cross-domain visibility they currently lack. They identify and remediate misuse on their own networks before it becomes a regulatory matter.	Infrastructure-layer harm reduces. Enforcement actions reduce. Industry effectively assumes the policing function previously impossible for either regulator or operator alone.
Stage 5	Industry-wide self-regulation reaches equilibrium. Regulatory enforcement normalises at a lower level as compliance improves across the ecosystem.	Ofcom is leaner, better-resourced, and focused on its highest-value work. Carriers are cleaner and commercially stronger. Consumers are better protected. The blind spot is closed.

when the backlog of visible but previously un-prosecutable violations is converted into enforcement actions. As compliance improves in response to the deterrent effect of consistent enforcement, the volume of violations reduces and enforcement revenue normalises downward. This is the intended outcome, not a failure of the model.

The long-term dividend is different in character. Once the industry has self-regulated to a significantly higher baseline of compliance, the ongoing cost of infrastructure-layer oversight falls dramatically, because the harm being overseen has reduced. Ofcom's enforcement resource, freed from the burden of manual retrospective investigation, can be permanently redeployed toward the online harm agenda and the next generation of digital governance challenges. The sector is cleaner, the regulator is more effective, and the consumer is better protected, at a lower ongoing cost than the current dysfunctional status quo.

The Compounding Effect - Why Every Year of Inaction Increases the Cost

The costs documented in this paper are not static. They are growing, and the rate of growth is accelerating. Understanding the compounding dynamic is essential to appreciating the urgency of the argument.

The AI Acceleration Factor
AI is not merely a feature of the threat landscape described in Paper One. It is the primary driver of cost escalation across every dimension documented in this paper. AI-generated voice fraud is growing at a rate that outpaces detection capability. AI-driven botnet infrastructure is scaling faster than manual security teams can track. AI-enabled phishing and smishing campaigns are operating at volumes that overwhelm the retrospective evidence cycles that current enforcement frameworks depend on. The 15 to 20 percent annual growth projection for telecom fraud losses is a conservative estimate that assumes no step-change in AI capability. In a market where AI capability is advancing at the rate observed since 2023, the true trajectory may be significantly steeper.

The Regulatory Gap Risk
As Paper Two identified, the UK faces a growing gap risk relative to comparable international frameworks. The EU's NIS2 Directive, in force from October 2024, introduces cross-domain incident reporting obligations that begin to address the fragmentation problem. The US CISA communications risk framework explicitly identifies AI-driven telecom threats as a top-tier national security concern. ETSI and 3GPP are developing cross-domain security standards that will eventually become market expectations. Each year that the UK does not develop equivalent frameworks, the cost of catching up increases, both in terms of the harm that accumulates during the gap period and in terms of the remediation required to bring UK infrastructure to international standards once the gap becomes untenable.

The Cost of Acting Under Compulsion
Regulatory reform under crisis conditions is consistently more expensive than reform undertaken proactively. When a sufficiently large or visible infrastructure-layer incident forces parliamentary attention, a scenario that is increasingly plausible given the trajectory of carrier breach events, the resulting regulatory response is likely to be faster, blunter, and more costly to implement than a considered, evidence-based transition to a unified intelligence model. The carriers that have invested in cross-domain compliance capability will be well-positioned; those that have not will face both the direct costs of rapid compliance and the reputational consequences of being seen to have delayed.

The cost of acting now is a fraction of the cost of acting after a crisis has forced the issue. The self-financing flywheel model makes the economics of proactive action compelling. The only argument for continued inaction is the belief that the crisis will not come, a belief that the incident record of the past two years does not support.

What Proportionate Investment in Prevention Would Cost

This paper has documented the costs of inaction across four stakeholder groups. It has argued that those costs are growing, compounding, and materially larger than headline breach figures suggest. It has presented a self-financing model for addressing them. The natural final question is: what would proportionate investment in the solution actually cost?

Paper Four in this series will set out the operational model for cross-domain resilience in detail. This paper confines itself to the high-level proportionality argument.

The investment required to build a shared infrastructure-layer intelligence platform, encompassing unified telemetry ingestion, cross-domain correlation, AI-aware detection, and evidence-grade reporting, is measurable and finite. It is not a programme of the scale associated with national infrastructure investment. It is a software and data architecture challenge of the kind that has been solved in comparable sectors, financial crime monitoring, environmental compliance, health and safety enforcement, at costs that are modest relative to the harms they address.

The proportionality case is straightforward. US$38.95 billion in annual global telecom fraud losses. US$10.22 million average breach cost per incident. Tens of billions of pounds in annual consumer fraud with a significant telecoms-enabled component. Enforcement revenue from infrastructure-layer violations currently going uncollected at scale. Against these numbers, the investment required to build the unified intelligence infrastructure

that closes the blind spot is not a financial question. It is a question of institutional will.

The Investment Case in One Sentence

The cost of building the shared intelligence infrastructure that addresses infrastructure-layer harm is a fraction of the annual cost of the harm it prevents, and the enforcement revenue it unlocks in its early stages is likely to exceed the cost of building it.

Conclusion - The Arithmetic of Inaction

This paper has attempted something that has not previously been done in the infrastructure-layer harm space: a complete accounting of what inaction costs, across all four dimensions that bear the burden of that cost. The arithmetic is not complicated. Carriers face direct breach costs, pipeline loss, customer churn, insurance premium resets, and reputational damage that collectively represent a multiple of the headline figures in standard breach analyses. Regulators face the resource drain of manual retrospective enforcement, the enforcement revenue they cannot collect because the evidence framework does not exist, and the institutional credibility cost of documented blind spots. Enterprise customers face security exposure, operational disruption, and AI-enabled fraud that operates invisibly inside trusted carrier infrastructure. Retail consumers face financial fraud, compromised home networks, downstream phishing, and the gradual erosion of trust in the communications infrastructure they depend on.

Add these costs together, apply the compounding effect of 15 to 20 percent annual AI-driven growth in fraud losses, factor in the regulatory gap risk from international frameworks that are moving ahead of UK equivalents, and the case for inaction becomes economically indefensible.

The self-financing flywheel model presented in Section 7 resolves the one objection that might otherwise give policymakers pause. The transition to unified infrastructure-layer oversight does not require new budget. It generates its own funding through the enforcement revenue it unlocks. It creates its own compliance incentives through the deterrent effect of visible, consistent, evidence-grade enforcement. And it reaches its own equilibrium, an industry that has self-regulated to a significantly higher baseline, overseen by a regulator that is leaner, more effective, and freed to focus on the challenges that genuinely require human judgment and political accountability.

Paper Four in this series will set out what the operational model for that transition looks like in practice. The economic and societal case for undertaking it is documented here. The question is no longer whether the cost of inaction is acceptable. It is not. The question is who leads the transition and when.

About This Series

This white paper is the second in a four-part series on cross-domain telecommunications resilience:

- **Paper 1** - Telecom Resilience in the Age of AI: A Cross-Domain Framework for National Communications Security (April 2026)
- **Paper 2** - The Regulatory Bandwidth Crisis: Why Infrastructure-Layer Harm Has Been Left Unseen (May 2026)
- **Paper 3** - The Economic and Societal Cost of Infrastructure-Layer Inaction (this paper)
- **Paper 4** - The Operational Model for Cross-Domain Resilience (forthcoming)

About the Author

Robert Lee is a senior telecommunications and digital infrastructure expert with over 25 years of experience across global carrier, cloud, and hyperscaler markets. He has held senior roles at Colt Technology Services, Verizon Business, Tata Communications, Sky Business, and Cogent/T-Mobile and more, managing global accounts and delivering complex cross-domain network solutions across Data, Voice, IP, subsea, data centre, and cloud environments. He is the founder of Panoptes Intelligence and the author of the Cross-Domain Resilience white paper series.

This briefing is part of the DFM Briefing Centre, providing in-depth analysis and practitioner insight on emerging digital, cyber, and investigative issues.

Digital Forensics Magazine (DFM) is an independent publication focused on digital forensics incident response, cyber investigations, and applied cyber security research.

DFM produces practitioner-led analysis, briefings, and technical features for professionals across law enforcement, industry, and academia.

Website: digitalforensicsmagazine.com LinkedIn: Digital Forensics Magazine