Bugs are normal.
The rest is a choice.

A silent blackhole on a shared peering fabric, the leaked customer traffic that came before it, and what an exchange owes the members and data subjects who aren’t aware of either

James A. T. Rice 12 June 2026

ON METHOD

A note on how this is told

- I’ll write this up using some of the principles of an AAIB or NTSB interim report.
- Those reports don’t name the airline. They name the aircraft, the registration, the sequence of events — and they look for cause, not blame.
- So this deck doesn’t name the exchange or the members involved. It also means it travels: anyone reading it later has the facts without the finger-pointing.
- Bugs are normal. The interesting questions usually aren’t about the bug.

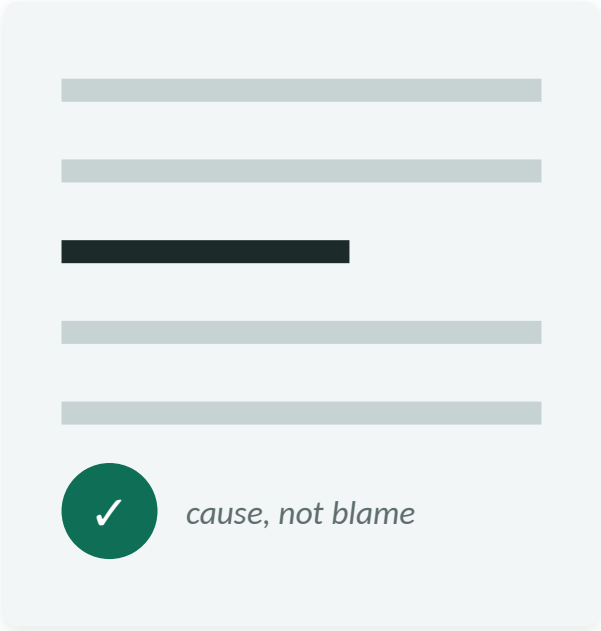

THE TECHNOLOGY

We never really left the nineties

The clock says 2026. The peering LAN presents much as it did then.

- On top, it presents as a shared layer-2 segment where members swap traffic.
- Underneath, layers accrete to keep that surface working: VXLAN, EVPN, proxy-ARP, proxy-ND and route servers.
- Control-plane signalling and data-plane traffic often no longer share a path — which is where blackholing becomes possible.
- Each layer added to keep the old surface working brings its own ways to fail.
- LAN hygiene needs active managing — otherwise the problems are needles in haystacks.

presents as

shared layer-2 segment (1990s)

underneath

VXLAN · EVPN

proxy-ARP · proxy-ND

route servers

others of the same kind — SPB-M, VPLS — turn up in earlier eras and other fabrics

MECHANISM

The mechanism, start to finish

A member changes a router or NIC — new MAC address

▼

Peers keep the old MAC in the neighbour cache

▼

Probes to the departed MAC are unknown unicast

▼

Delivered to the Juniper proxy-ND — it answers from any MAC, with no target-link-layer address

▼

RFC 4861 §7.2.5: that confirms reachability without correcting the MAC. The entry stays REACHABLE

▼

Traffic flows to a MAC that left — silently blackholed, while route-server BGP stays up, wrongly signalling reachability

In plain terms

A neighbour entry should expire and re-resolve once the MAC behind it changes. Here it never does: the exchange's proxy keeps answering probes for the address that left, so the cache stays confident in a dead MAC.

Everything sent to that neighbour is delivered nowhere — while the layers above it, the route-server BGP session included, still look healthy.

2021–Jan 2025 the replies came back via the flooded path; since then, via the proxy-ND. The effect is identical — perpetuated incorrect neighbour entries.

MECHANISM

You can watch it on the wire

The kernel probes the stale entry. The reply comes from the new MAC, with no correction:

```
NS  who-has 2001:db8::a05:1
    tgt lladdr 02:00:00:00:0b:fb   (old MAC)

NA  from       02:00:00:00:0b:fc   (new MAC)
    length 24,  no target lladdr
```

the entry stays

REACHABLE

pointing at the old MAC

The NA confirms reachability (§7.2.5) but carries no link-layer address, so the kernel never corrects the MAC.

THE SYMPTOMS

One root cause, two safety problems

Confidentiality

2021 – Jan 2025

Mis-forwarded to hundreds of member ports. Web domains visited by users were exposed through HTTPS SNI; passwords, full URLs and email contents travelled in cleartext. Any of those receivers could log, retain or index it.

resolved 7 January 2025

Availability

now

Mis-addressed traffic is silently dropped. Confirmed November 2024 by a downstream FTTP operator: “very” bad for safety of life (regarding 999 calls subject to being affected).

unresolved for years

THE SCALE

How many people?

Even at my own port, one instant in November 2024 caught traffic to and from networks including:

Apple	**Google**	**Meta**	**Microsoft**
Netflix	**Cloudflare**	**BBC**	**Akamai**

… along with traffic to one of the two largest eyeball networks in the UK.

Between them, these networks carry the everyday traffic of most of the country. Sustained over nearly four years, the communications exposed plausibly amount to those of up to tens of millions of people.

TESTING & REMEDIATION

What was fixed, and what wasn’t

On 7 January 2025 the exchange deployed a filter dropping unknown unicast on the affected switches. That ended the flooding — the disclosure. The blackhole, years on, remained unresolved.

How it affects members depends on how they peer:

- Bilateral-only peers see a BGP session drop — visible, no blackholing.
- Route-server-only peers see nothing. Traffic is silently blackholed; BGP stays up.
- Peers using both: the bilateral drops, but route-server-learnt traffic is still blackholed.

Relayed from Juniper JTAC (May 2026): IPv4 frames to an unknown destination MAC are suppressed as expected, on both local and remote interfaces. Repeating the test with IPv6 neighbour solicitations, both local and remote flooding still occur — the filter reduces the count but does not stop it. JTAC attributes this to some other, as-yet-unidentified mechanism, now under investigation; a case is open with product management for a filter that would catch it.

LATEST

Progress, the morning before this talk

- On 12 June 2026 the exchange deployed a change on the switches carrying the affected members.
- Another affected member reports their stale entries now resolve correctly.
- I have not yet confirmed it for my own connection — so this is a likely fix, not a double-verified one.

After years of recurrence, this is real movement on the technical fault. The questions the rest of this talk raises — about notice, and about decisions — are unchanged.

This is where it stops being a networking talk.

The rest is about how we look at failure.

FAILURE ANALYSIS

Reality cannot be fooled

> *For a successful technology, reality must take precedence over public relations, for nature cannot be fooled.*

Richard Feynman, Rogers Commission report on Challenger, 1986

The O-ring was the cause. The story was an organisation that had quietly got used to its own warning signs. Good failure analysis keeps looking past the broken part to the system that let it break.

FAILURE ANALYSIS

The same incident, as a graph

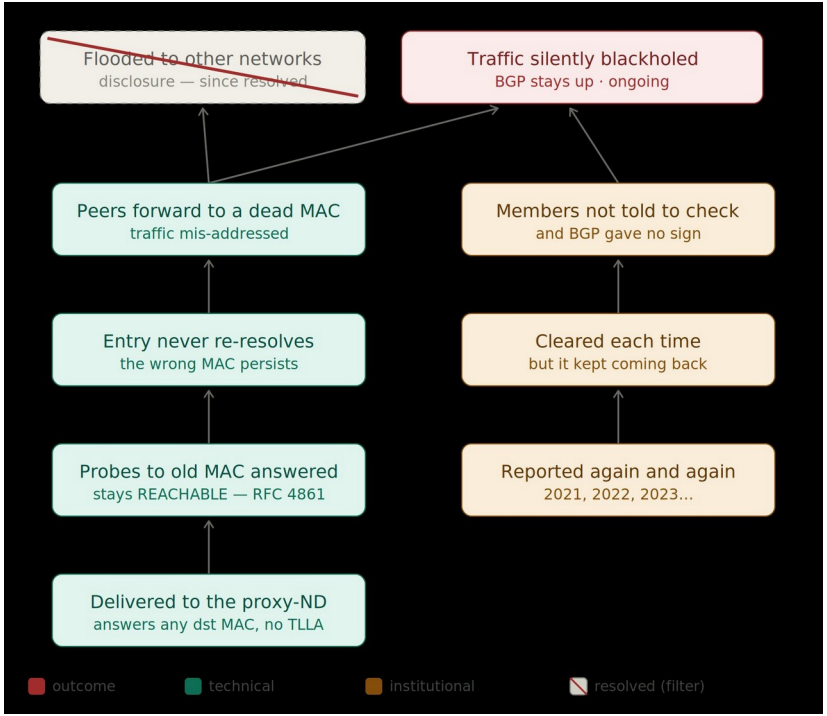

Two chains, doing two different jobs.

The technical chain makes the blackhole happen: the proxy-ND keeps refreshing a dead entry. The institutional chain makes it persist, quietly — reported for years, cleared each time but recurring, and never flagged to the membership.

Flooding to other networks, the disclosure, is struck through. That part was fixed. The blackhole wasn't.

THE COMMUNICATION GAP

None of this would shock an operator

- Engineers understand bugs.
- Large layer-2 fabrics are genuinely complicated.
- Vendor bugs happen, and JTAC cases take months.
- None of that would surprise anyone who runs a network.

What would concern many members is the communication, not the bug.

Being told is what lets each member decide whether, and how, to mitigate.

ACCOUNTABILITY

Acceptance comes before confidence

Unless an organisation has accepted its mistakes, we have little reason to be confident it won't repeat them.

> *"As we have not been able to reach the other two networks for a resolution we will shut down their ports while we work on a resolution with them."*
>
> a senior executive, by email · 15 November 2024

> *"We will be more careful to avoid stating proposed actions to members until the course of action is fully agreed and irrevocable."*
>
> a senior executive · June 2025

In all this correspondence, this is the only specific misstep the exchange has acknowledged.

The commitment it made was to make fewer commitments.

ACCOUNTABILITY

Decisions the exchange made unilaterally

- 15 November 2024: the exchange committed to shutting down the remaining misbehaving member ports that day, which would stop the traffic disclosure and disruption. With no discussion or notice, that commitment wasn’t implemented.
- So the leak and disruption continued until the unknown-unicast flooding was filtered, on 7 January 2025.

Leaving a mis-forwarding member connected is a decision that affects every other member and their customers. It was a decision made for them, and without consulting them or telling them.

THE EXCHANGE'S POSITION

What the exchange says

- It reported itself to the ICO, which did not open an investigation.
- But it would not show members what it sent, or the ICO's reply — despite telling the membership on 18 November 2024 it would share the outcome as far as it could.
- It is confident its conduct was entirely lawful.
- Its preliminary view is that there was no reportable incident under UK GDPR or NIS.
- It considers the matter closed.

> *"We are confident that our conduct was entirely lawful and we consider the matter closed."* — *a senior executive, 2025*

That may be right. This talk is about the part it leaves unanswered — and the people it affects.

“

The rights of the data subjects whose communications were forwarded outside their intended path, and disrupted.

They are not present in this correspondence, are largely unaware of what happened, and have no voice in it. They are the people whose rights and safety the statutory frameworks exist to protect.

from my own correspondence, May 2026

THE QUESTION

“

The question I’d most like answered remains whether affected members must notify their own customers that their communications were disclosed to other networks and disrupted — and, importantly, the reasoning, since a member can only discharge that duty if they understand the basis for it.

The frameworks in play

The statutes a fault like this engages.

Communications Act 2003, ss.105A–105K (Ofcom). A "security compromise" covers the confidentiality of signals conveyed (s.105A(2)(c)) and availability or functionality (s.105A(2)(a)) — so disclosure and blackhole both appear to qualify. Duties to act (s.105C), report to Ofcom (s.105K) and inform users (s.105J).

NIS Regulations 2018 (Ofcom). An IXP operator is a deemed operator of an essential service at ≥50% UK IXP market share (by interconnected ASNs), or ≥50% of global routes: reg 10 security, reg 11 notification within 72 hours.

UK GDPR Arts 33–34 / DPA 2018 (ICO). Notify the ICO; tell individuals where the risk is high; document even a non-reportable breach (Art 33(5)).

PECR 2003, reg 5A (ICO). Every personal data breach notifiable, with no materiality threshold (within 72 hours, raised from 24 in 2025); notify the subscriber or user.

Electronic Communications (Security Measures) Regulations 2022, reg 7 (Ofcom). Manage third-party-supplier security risk, with flow-down and contingency.

OPEN QUESTIONS

And which obligations applied?

What did the exchange conclude? Asked in writing on 11 May; still open, posed here for response.

- Manage and remedy security compromises — the disclosure and the silent loss of traffic both appear to qualify.
- Report serious incidents to the regulator, as critical national infrastructure.
- Assess and report a personal data breach — and tell the people affected.
- Notify every breach — communications providers get no “too small to matter” threshold.
- Manage the supplier risk passed down the chain — impossible for a fault a provider was never told about.

An explicit position with a brief reason on each — even a reasoned “does not apply” — can be discussed here. Silence cannot.

THE PRINCIPLE

On optics

For notifying members about potential safety matters, optics should never be part of the equation.

And if that is hard to accept, consider that the optics are always far worse if the exchange makes decisions not to notify members — decisions that harm members' correct traffic delivery and the safety of persons and data subjects.

my position to the exchange, 2026

Is this fine?

We can fix the bug. The open questions are whether the people whose communications were exposed are ever told — and whether the exchange makes better decisions next time.

Thank you. James A. T. Rice · Jump Networks

This material was also presented at RIPE 92 (20 May 2026) and NetLdn 73 (10 June 2026):
https://ripe92.ripe.net/programme/meeting-plan/sessions/84/YZVWHG/
https://netldn.uk/2026/06/04/netldn-73-10-06-2026-date-change/