May 2025
Geoff Huston

Analysis of a Route Leak

The interdomain routing fabric of the Internet is a somewhat chaotic space. It is not a "centrally managed" space by any means. There is no single entity that is responsible for the Internet's routing environment. No permissions are required to join, and no notice required to leave. A network doesn't need to meet any pre-determined qualifications or demonstrate any level of operational competence in operating a local routing instance. You select one or more routing neighbours, set up BGP peering sessions with them (and likely make a payment or two), and if do it correctly your routes will propagate across the Internet, and you will also receive every other network's routes.

Such arrangements, or the lack thereof, seems just too informal, too anarchic and too unstable to form the foundation of the world's communication system. Yet here we are. And, surprisingly, it works! Well, it mostly works most of the time. And from time-to-time anomalous things happen. In this article I'd like to analyse just one of the instances when things didn't go according to plan, showing how the routing anomaly was visible, and how it could be mitigated.

The Leak

Figure 1 shows the total size of the BGP routing table, as seen by AS131072, using hourly snapshots across April and May 2025.

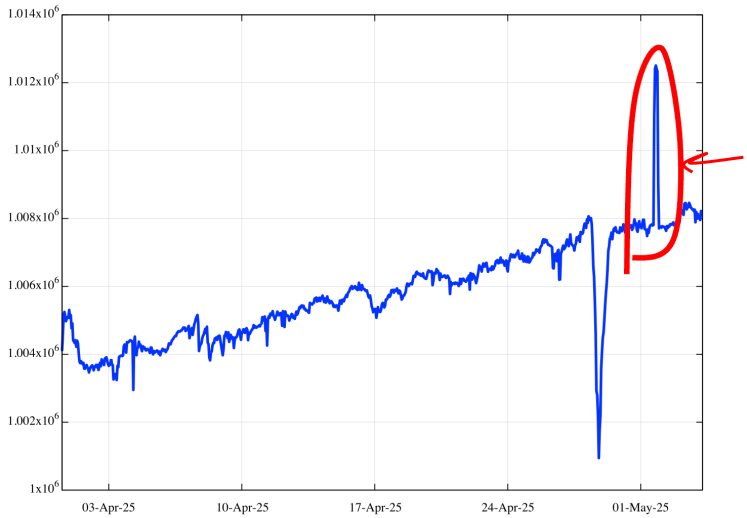

Figure 1 – BGP IPv4 Routing Table April – May 2025

I'd like to look at that spike in the total route count that occurred on the 1st May in further detail. Between 16:00 UTC on the 1st May and 18:00 UTC on the same day, the routing table grew by some 4,500 routes. They were visible for the next four hours, and then disappeared, as shown in further detail in Figure 2.

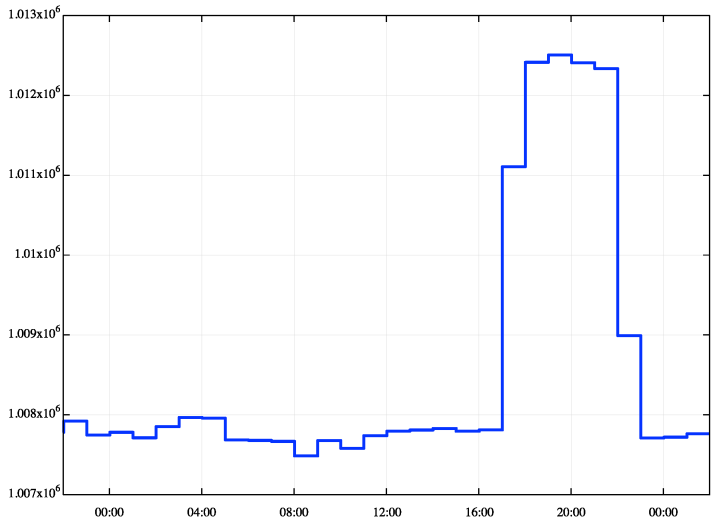

Figure 2 – BGP IPv4 Routing Table, 1 May 2025

We can provide more detail on this event by looking at the log of BGP updates that were received over this period. In Table 1 I've used the update log to aggregate the changes to the routing table in 5-minute buckets. It appears that this route leak occurred in two stages. The first was a leak of 3,365 routes between 16:45 and 16:50, followed by a second leak of 1,141 routes an hour later between 17:50 and 17:55 (Table 1). The leaked routes were withdrawn some three hours later, again apparently in two stages, one hour apart (Table 2).

These times are local times when the new advertisements were received by local routing instance. A number of BGP implementations use a Minimum Routing Advertisement Interval (MRAI) time, which generally uses a random interval between 27 and 30 seconds, so the original leak may have occurred a couple of minutes before the time when the updates were received by my local BGP instance.

Time	Size	Change
16:25	1,001,322	12
16:30	1,001,272	-50
16:35	1,001,270	-2
16:40	1,001,255	-15
16:45	1,001,297	42
16:50	1,004,662	3,365
16:55	1,004,615	-47
17:00	1,004,586	-29
17:05	1,004,597	11
17:10	1,004,608	11
17:15	1,004,591	-17
17:20	1,004,640	49
17:25	1,004,634	-6
17:30	1,004,663	29
17:35	1,004,719	56
17:40	1,004,732	13
17:45	1,004,757	25
17:50	1,004,743	-14
17:55	1,005,884	1,141
18:00	1,005,890	6
18:05	1,005,993	103

Table 1 – Route Leak

Time	Size	Change
21:10	1,005,814	-21
21:15	1,005,838	24
21:20	1,005,817	(21)
21:25	1,005,840	46
21:30	1,005,593	-247
21:35	1,005,211	-382
21:40	1,004,740	-471
21:45	1,002,366	-2,374
21:50	1,002,375	9
21:55	1,002,386	11
22:00	1,002,392	6
22:05	1,002,438	46
22:10	1,002,410	-28
22:15	1,002,445	35
22:20	1,002,409	-36
22:25	1,002,420	11
22:30	1,002,205	-215
22:35	1,002,188	-17
22:40	1,001,183	-1,005
22:45	1,001,184	1
22:50	1,001,212	1

Table 2 – Leak Recovery

Who Leaked?

The BGP update log contains a copy of every received BGP Update message. For an Update the message contains the IP address prefix being updated, and the AS Path. For Withdrawals, it's just the IP address prefix. This information allows us to track the total number of advertised prefixes

per originating AS. Every time we see a prefix being advertised with a new originating AS we will add to the count of prefixes for this, and similarly we'll process explicit and implicit withdrawal messages.

During the period when these routes were announced AS 22773, a network operated by Cox Communications, a major US retail ISP, announced some 4,651 additional routes. Some three hours later the same AS withdrew 4,662 routes. A timeseries plot of the total number routes originated by this network is shown in Figure 3.

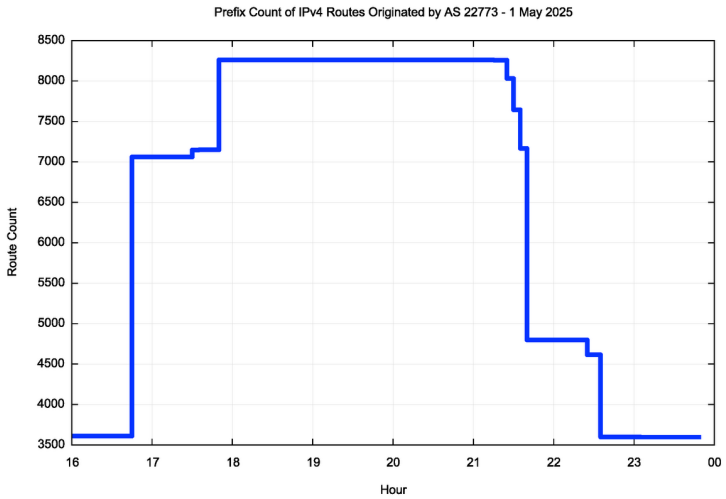

Figure 2 – BGP IPv4 Prefixe Count for AS22773, 1 May 2025

Who Saw this Leak?

A leak of some 4,600 routes is not totally trivial. While it represents just 0.5% of the total IPv4 route collection, it's still far higher than the "normal" level of BGP activity and would normally attract some mention in routing circles.

However, this particular event has not been noticed in many parts of the network. Surely the role of BGP is to reliably flood all routing information, good or bad, to all part of the network, and if that's the case then every BGP speaker should've seen this route leak. But this was not the case here.

The explanation lies in the use of RPKI-based routing security tools. The network I'm using as my observation point deliberately does not perform any testing for received routes for RPKI-based validation, as I wish to observe both RPKI-valid and RPKI-invalid routes in order to assemble a picture of the levels of adoption of RPKI validating in the interdomain routing space.

When I look at the set of 4,651 routes leaked from AS22773 on 1 May 2025, most of these routes (4,644 routes) would be marked as RPKI-invalid by a RPKI-aware BGP speaker, and the local BGP instance would not have learned these routes. Just 7 of the leaked routes were not covered by a valid ROA.

Cox has a pretty complete set of published ROAs for the address prefixes that it advertises. The network currently publishes 6,021 ROAs, which covers 95.06% of its advertised IPv4 address prefix set (Figure 4).

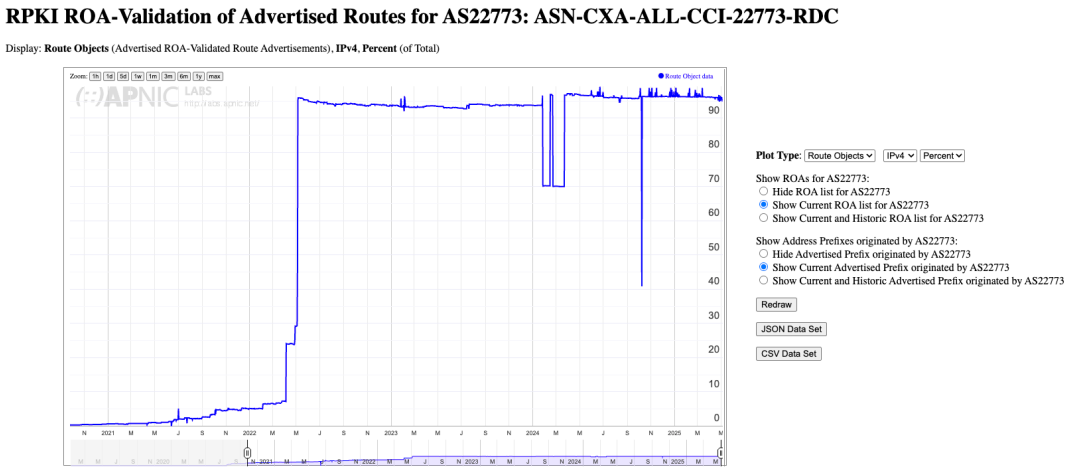

Figure 4 – ROAs published by AS22773 – from APNIC Labs Report

The network also appears to have enab led the dropping of learned routes if they are RPKI-invalid since early July 2025 (Figure 5). So only those networks not performing RPKI-based route filtering would've seen this route leak.

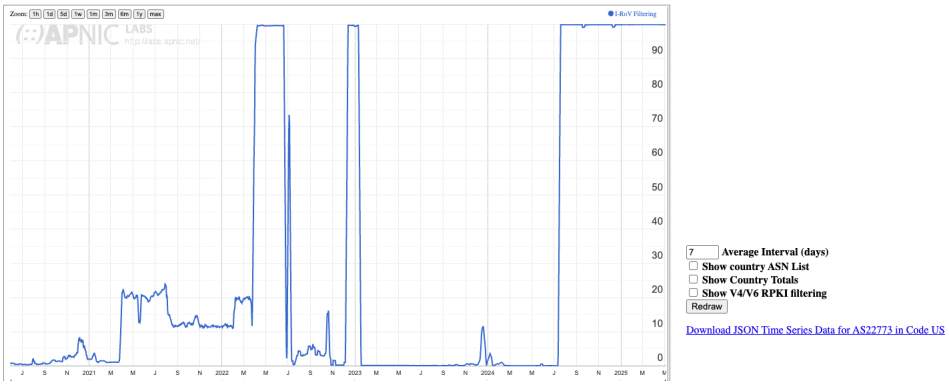

Figure 5 – RPKI Drop Invalid rate for AS22773 – from APNIC Labs Report

Without widespread use of RPKI-based AS Path validation (which some would say is coming soon, and others would claim that its protracted gestation period in the IETF implies that there is little interest in its adoption, ever!), these RPKI filters relating to route origination will not foil a determined attacker. However, as this case study shows, the careful use of published ROAs that precisely match the intentions of the network in advertising prefixes, can be used by others to filter out inadvertent route leaks. If may be a somewhat heavyweight approach to route leak detection and filtering, but as long as networks maintain an accurate collection of ROAs that match their routing intentions, this can be a very useful response to routing leaks.

The one question left in my mind is: Why doesn't AS22773 apply the same RPKI filter to its outbound advertisements? If it had done so, then the routes would not have leaked out at all!

You may have also noticed an earlier event in Figure 1, where the number of routes dropped by some 5,500 entries for a period of 24 hours from midday on the 28th April.

This is not a RPKI and route leak issue, as withdrawals cannot be validated by RPKI.

The cause of this is easier to identify, and this routing anomaly can be attributed to the widespread blackout that occurred in the Iberian peninsula on that date, where power was lost across large areas of Spain and Portugal.

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

Author

Geoff Huston AM, M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net