

March 2026
Geoff Huston

The Why and What of the CIDR Report

For some time, I have been looking after a routing analysis report called the "[CIDR Report](#)". Here I'd like to explain the reasons for this report, and what is in the report and share some thoughts as to its usefulness today to the Internet routing community.

Why the CIDR Report

To place this into context we probably need to head back to the Internet as it was in the 1980's. The milestone transition of the Internet from a small-scale US ARPA-funded research project investigating wide-area packet switching technologies into a technology that is the core platform of the entire global digital communications environment happened in the middle of the 1980's. That was the time of a decision by the US National Science Foundation to act as the lead funding agency for a national research backbone computer network, connecting US academic and research institutions to national supercomputer centres and each other, via a two-tier structure of eleven regional networks and the NSF-supported NSFNET backbone, all using the IP protocol. It wasn't the only national research and scientific backbone network, as at the time NASA operated the NASA Science Internet (NSI) and the US Department of Energy supported a High Energy Physics Network (HEPNET). While these latter two networks were originally based on the DECNET network protocol, they were upgraded to use multi-protocol routers to add IP support around that time. The result was a quickly growing internetworking routing problem, and the early routing tool, namely the Gateway-to-Gateway Protocol (GGP), was struggling with growing pains.

In January 1989 IBM's Yakov Rekhter and Cisco's Kirk Lougheed came up with the [Border Gateway Protocol](#) (BGP). It was a simple distance vector protocol with four critical aspects:

- It used TCP transport, so once a BGP speaker had informed a BGP peer of an update it could assume that the neighbour then "knew" this information and there was no need to re-send the same information it at a later time within the same BGP TCP session.
- BGP attached a path vector to each route update to prevent the classic distance vector count-to-infinity behaviour to detect routing loops.
- BGP was positioned as a dedicated inter-domain routing protocol and had no role in internal topology management. Each domain, or network, could independently choose a routing protocol to maintain its internal topology. BGP did not attempt to address everyone's routing needs!
- BGP supported an inter-network of peer networks, and there is no entity or role that controlled or orchestrated the overall inter-domain space.

The protocol design of BGP struck a resonant note with the emerging Internet. Each component network was able to operate their own network in a largely autonomous manner, including the choice of internal routing protocol, and the collection of inter-network routing policies used in connecting to other networks (via BGP, of course) was for each constituent network to determine and implement. In this

sense they were "autonomous" networks. This self-organising structure of peer networks in BGP was a very good match to the needs of the emerging Internet, and this has continued to this day.

Addresses – From Class to Classless

A routing protocol passes reachability information about endpoint addresses, and the internal structure of the network's address plan has a direct bearing on the architecture of the routing protocol.

The Internet's address plan was similar to the E.164 international numbering plan for the global telephone network, where each component network was assigned a unique address block with a unique prefix value, and individual connected devices had a local device address which was unique within that component network. The difference between these two address plans lay in the definition of these *component networks*. In the E.164 telephone address plan, each component network was a national telephone system, so there was a "natural" limit of some 200 such network prefixes correspond to the number of national telephone networks. The Internet started from a different point, where component networks were interconnected local area networks. If an entity, such as a university or research institute, operated a number of distinct local area networks, then they used a number of distinct network prefixes.

A more significant point of difference was IP's choice of a stateless packet forwarding paradigm, eschewing the support of dynamic virtual circuits. This allowed for far simpler (and cheaper) networks (as they did not incur the overheads of supporting the operation of virtual circuits as overlays on the underlying network) but required that every packet had its destination address inscribed into the IP packet header. The strictures of packet forwarding capacity came down on the side of a fixed length IP address that contained both the network prefix and the local device address. This presented the network's architects with an almost intractable problem: how to partition a fixed sized address into network prefix and device address parts so that the diversity of local networks (large and small) could be encompassed.

Prior to 1994 the address plan used by the Internet was what we called *class-based* addresses, and the 32-bit address pool was divided into three sizes of address blocks. There were 126 distinct *Class A* 8-bit network prefixes, each able to encompass 16,777,216 individual device addresses, 16,383 *Class B* 16-bit network prefixes, each with 65,536 device addresses, and 2,097,152 *Class C* 24-bit network prefixes, each with 256 device addresses.

The early days of the expansion of the Internet into the research community in the late 1980's saw rapid consumption of Class B addresses, to the point where exhaustion of Class B addresses was an imminent proposition by the end of that decade. We needed to shift away from a static class-based address plan within the Internet, including changing the routing protocols that we were using.

Scaling Routing

The second issue was that the number of routes (network prefixes) being passed into the inter-domain routing space was rising rapidly. The stateless datagram IP architecture meant that every BGP router that connected to the inter-domain space was carrying a full routing table of routes in its forwarding memory structures.

As the number of routes increased, this placed pressure on the amount of memory used by routers to hold these routes, and the amount of time that was consumed in performing an address lookup within this structure. At the same time increasing circuit speeds in wide area networks was decreasing the available per-packet processing time for routers, so this was a classic case of pressure to perform more work in less available time.

A view of the size of the Internet's routing tables up to the end of 1994 is shown in Figure 1. The period from 1988 to January 1994 uses monthly routing reports from Merit, the operator of the NSFNET backbone, recording the size of the BGP routing table in the NSFNET. From January the data series shifts to an hourly collection. This 6-year view encompasses a significant transition for the Internet in March 1994. In the early 1990's the chosen response for the Internet was to dispense with the Class-

structure in IP addresses, and early 1994 a new version of BGP (BGP-4) was deployed across the Internet that supported "Classless Internet-Domain Routing" or "CIDR".

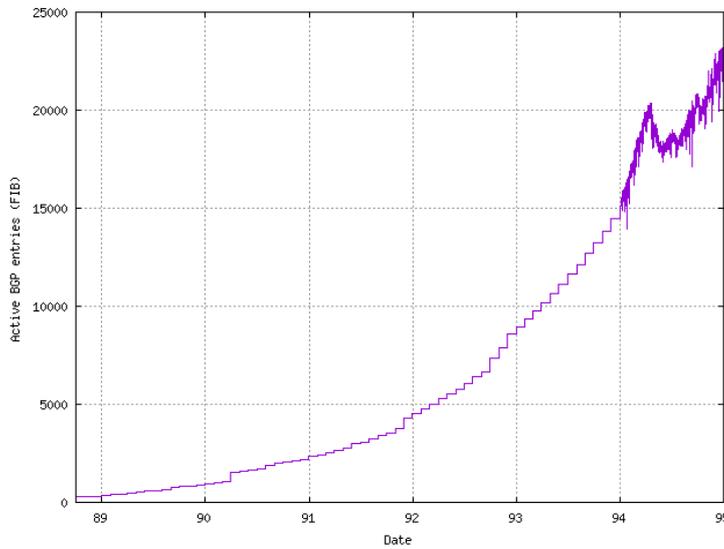


Figure 1 – Size of the Internet's Routing Table 1988 - 1995

The period from January to March 1994 saw a 25% growth in the size of the routing table and the introduction of CIDR in around March 1994 saw the size of the routing table drop from 20,000 entries to 18,000 over a couple of weeks.

Over the ensuing five years the routing table continued to grow, but the growth rate of the routing table was far lower than the peak growth rate observed the start of 1994. The growth trend from 1995 until 1999 was a largely linear trend, growing by some 8,000 additional routing entries per year. This slower growth rate in the routing table could be accommodated within the operational lifecycles of core routing equipment, so up until the start of 1999, the growth rate of the routing system was not a cause for significant concern (Figure 2).

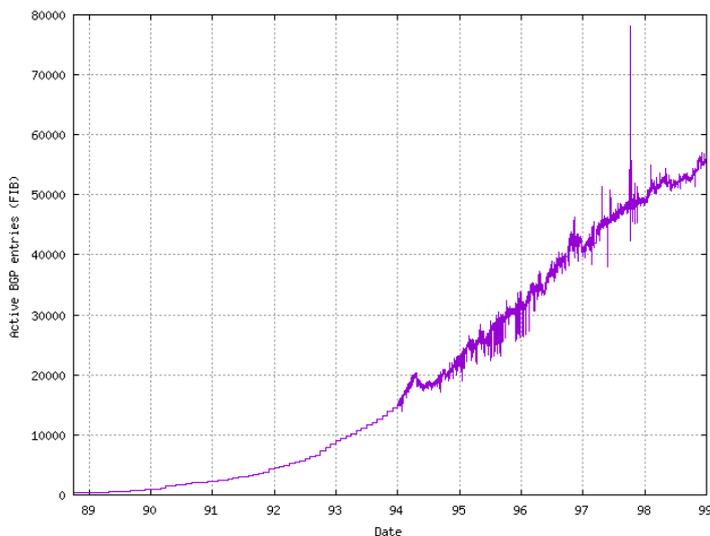


Figure 2 – Size of the Internet's Routing Table 1988 – 1999

The picture changed once more at the end of the 1990's. The two-year period across 1999 and 2000 was the peak of an Internet boom, and the routing table had resumed an exponential growth trend. This was a short-lived period of Internet euphoria, and the ensuing bust appeared in the routing table in 2001, where the routing table holding steady at some 105,000 entries across all of 2001. This period of stasis

was short-lived, however, and routing table growth resumed in 2002, and once more the growth trend was exponential. The drivers at this time were the adoption of DSL "always on" residential services in the consumer markets of many economies at the start of that decade, replacing the dial-up analogue modem as the dominant connection method for Internet services. Later in that decade the Apple iPhone was launched, which heralded the expansion of the Internet into the realm of mobile services. In the final three years of this decade the growth of the Internet's routing table was 100,000 additional routing table entries, a growth rate of 50% from the 200,000 entries at the start of 2007.

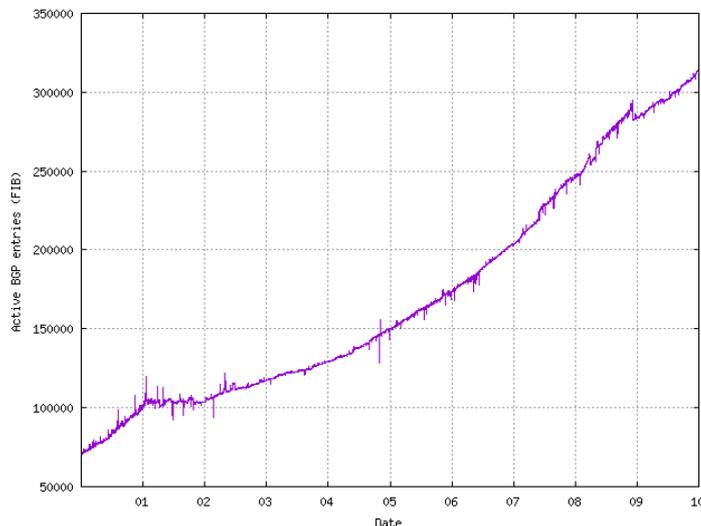


Figure 3 – Size of the Internet's Routing Table 2000 – 2009

More Specific Routes

What was driving this growth in the routing system? The explosive introduction of mobile devices into the Internet was one of the driving factors, but this growth pressure was moderated by the increasing use of network address translation (NAT) in these mobile access networks. Another factor was also apparent, namely the use of *more specific route advertisements*. This is a routing practice that is enabled by CIDR, namely the use of overlapping route advertisements. For example, a route for 192.168.2.0/24 is a *more specific* route of a covering aggregate of 192.168.0.0/16.

Why would network operators do this? There is a useful behaviour of the BGP path selection algorithm where BGP will select a more specific route in favour of a covering aggregate route. For example, a network with external connections to providers A and B may want to balance the incoming traffic across the two connections. One way of achieving this is for the local network to advertise the covering aggregate across both connections but then augment this with more specific route advertisements to the provider where you are wanting to increase the incoming traffic volume. Where a network has a rich set of external connections and wants to optimise the incoming traffic volume across these connections, the use of more specifics can be a very effective means of *traffic engineering*.

There is another motivation in the area of *defensive routing*. An attacker may inject more specific routes for a target's address prefixes into the routing system in an effort to divert traffic, and thereby steal all the traffic being directed to these addresses, by virtue of BGP's preference to use more specifics over aggregates. By defensively advertising more specific routes, the potential damage radius of a hostile more specific route can be minimised.

We can look at the total number of more specifics in the BGP routing table across the same decade in Figure 4.

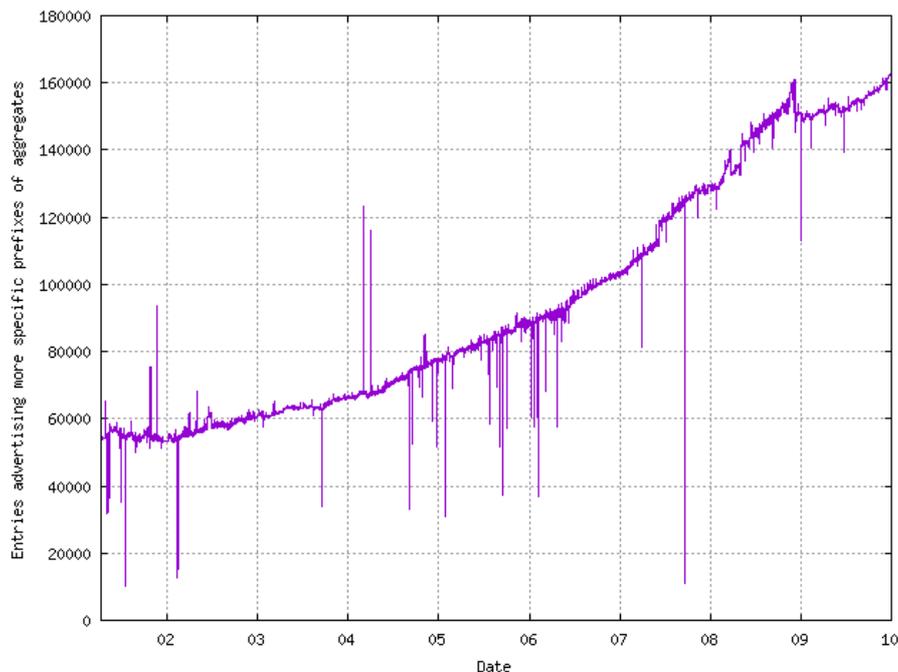


Figure 4 – Count of More specifics in the Internet's Routing Table 2000 – 2009

In the year 2000 more specifics accounted for some 55% of the total count of routes in the Internet's routing table, yet they encompassed less than 10% of the advertised address span. If this trend were to continue unabated, the routing table would quickly grow to sizes that exceeded the capabilities of routing hardware then available to network operators, which would be a crisis point for the Internet.

This is a classic form of "*The Tragedy of the Commons*", where individual self-interest for a network operator lies in exercising greater control over traffic flows and greater levels of resilience against hostile routing attack through control of advertisement of more specific routes into the inter-domain routing system, but the collective outcome of these actions results in a bloating of the inter-domain routing space, leading to an Internet that is large enough to be simply unrouteable.

There are few ways that we can exercise control over this form of behaviour by network operators. Don't forget that one of the essential attributes of the BGP design is that there is no overall control function and no one is in control. It's a network of peer networks. Collective action by hundreds of thousands of individual network operators to limit the extent of advertisement of more specific routes across the global Internet is not an available option.

One possible approach lay in the approach of increasing the level of public awareness of network operators' routing practices, illustrating their network's contribution to the current state of the Internet's routing table. This was the motivation for the publication of the CIDR report. If we can't directly control the use of more specific advertisements in the inter-domain routing space, then we can expose the extent of this behaviour and name those network operators that are more profligate in advertising these more specific routes. The hope here is that some form of self-moderation or peer pressure might influence these network operators to reduce the impact of their network's advertisements within the larger picture of the Internet's inter-domain routing environment.

And that's the rationale for the CIDR Report.

Let's now move on to the report itself.

What is in the CIDR Report

The CIDR Report was introduced in the late 1990's. As we've noted, the intention of this report was to expose to a broader level of visibility the impact of the advertising of more specific routes into the global

BGP routing system, and identify those networks whose routing practices were have a significant impact on the bloat that was evident in the Internet's inter-domain routing system.

The CIDR Report was first produced by Tony Bates, then taken on by Philip Smith, who then passed the baton to me. These days the report operates on a platform provided by APNIC, using snapshots of the global routing system assembled by BGP at AS 131072, a network number used by APNIC Labs. The report is produced on a daily basis and is available at <https://www.cidr-report.org>.

A snapshot of this header of the report is shown in Figure 5.

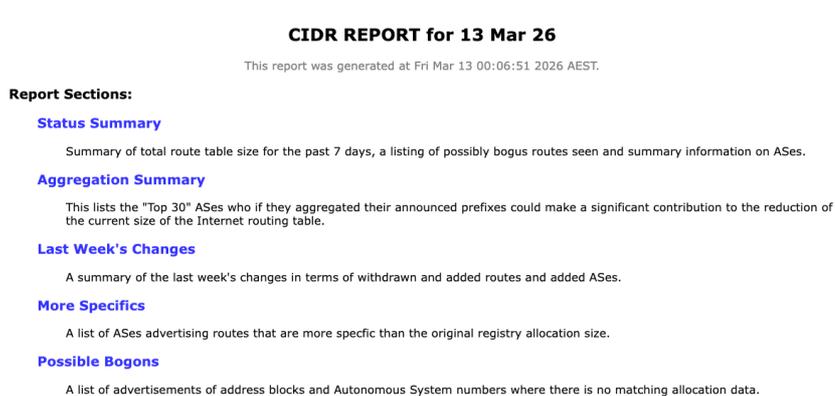


Figure 5 – Cidr Report Header

The report has five parts, a status summary, an aggregation summary, the changes over the past week, the top 20 networks advertising more specific routing entries, and a list of so-called "Bogon" routes.

Status Summary

The status summary is a condensed summary of the state of aggregation in the BGP routing table. The table shows the size of the BGP IPv4 routing table, and the size if all redundant more specific were removed, for each day over the past week, and a plot of the hour-by-hour total size of the routing table for that week. (Figure 6). (For those who were wondering, yes, there is an IPv6 version of the CIDR Report, containing the same information for the IPv6 interdomain environment, available at <https://www.cidr-report.org/v6>).

Table History

Date	Prefixes	CIDR Aggregated
06-03-26	1060364	595653
07-03-26	1060369	596247
08-03-26	1060586	596608
09-03-26	1060803	597331
10-03-26	1061105	598175
11-03-26	1060933	598256
12-03-26	1061773	598926
13-03-26	1061589	599673

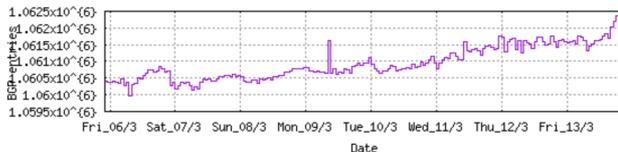


Figure 6 – CIDR Status Summary

A *redundant more-specific* route is a prefix advertisement that has an identical AS Path to that of its immediately encompassing aggregate route, so in terms of the application of local routing policies the more specific route does not have any tangible impact to the local forwarding function. In March 2026 some 461,596 advertised routes fall into this class of more specific routes.

This initial section of the report also reports on some statistics relating to the number of Autonomous System Numbers (ASNs) in the routing table (Figure 7).

AS Summary

78361 Number of ASes in routing system
 27121 Number of ASes announcing only one prefix
 14456 Largest number of prefixes announced by an AS
 AS16509: AMAZON-02 - Amazon.com, Inc., US
 227731456 Largest address span announced by an AS (/32s)
 AS749: DNIC-AS-00749 - United States Department of Defense DoD, US

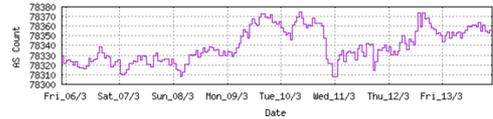


Figure 7 – AS Summary

Aggregation Summary

The next section of this report lists the 30 networks which have the highest count of these redundant more specific route advertisements (Figure 8).

Aggregation Summary

The algorithm used in this report proposes aggregation only when there is a precise match using AS path so as to preserve traffic transit policies. Aggregation is also proposed across non-advertised address space ('holes').

--- 19Mar26 ---

ASnum	NetsNow	NetsAggr	NetGain	% Gain	Description
Table	1062741	602294	460447	43.3%	Gain
AS9808	13836	3153	10683	77.2%	Gain
AS8151	12856	5109	7747	60.3%	Gain
AS12479	7201	288	6913	96.0%	Gain
AS7545	5959	544	5415	90.9%	Gain
AS39891	4969	97	4872	98.0%	Gain
AS9498	4875	342	4533	93.0%	Gain
AS16509	14479	10364	4115	28.4%	Gain
AS18403	4554	557	3997	87.8%	Gain
AS7155	3861	89	3772	97.7%	Gain
AS22773	4779	1076	3703	77.5%	Gain
AS4538	5155	1923	3232	62.7%	Gain
AS6327	3709	502	3207	86.5%	Gain
AS7713	3676	534	3142	85.5%	Gain
AS56046	3476	647	2829	81.4%	Gain
AS396982	3291	499	2792	84.8%	Gain
AS7552	4031	1328	2703	67.1%	Gain
AS45899	3345	829	2516	75.2%	Gain
AS24560	3014	683	2331	77.3%	Gain
AS17561	2927	645	2282	78.0%	Gain
AS12389	3198	1086	2112	66.0%	Gain
AS174	3274	1273	2001	61.1%	Gain
AS11492	4646	2709	1937	41.7%	Gain
AS45090	2190	396	1794	81.9%	Gain
AS11830	1953	444	1509	77.3%	Gain
AS45609	2019	515	1504	74.5%	Gain
AS7029	2580	1103	1477	57.2%	Gain
AS4755	2237	785	1452	64.9%	Gain
AS398781	1707	363	1344	78.7%	Gain
AS20057	1456	113	1343	92.2%	Gain
AS13999	2458	1126	1332	54.2%	Gain
Total	137711	39122	98589	71.6%	Gain

Report: [Aggregation Top 20 summaries for last 7 days](#)
 Report: [Full Aggregation Details](#)

Figure 8 – Top 30 Aggregation Opportunities

Each AS in this list points to an individual report that shows in detail the suggestions in terms of route withdrawals that could minimise the routing footprint of this network, while preserving the intended functionality of the network's traffic engineering policies.

Last Week's Changes

The next section looks at the changes to the routing table that occurred in the past 7 days. This is a section with multiple lists.

The first list is a list of those AS numbers that originate routes which were not visible in the routing table 7 days ago, and the number of routes they are currently originating. The second list is those AS numbers which have an increased total number of originating routes over the past week, ordered by the number of additional routes. The next list is those AS numbers which have reduced the total number of originating routes over the past week, ordered by the number of decreased routes. The next two lists are the ordered of networks which have added and removed routes in the past week.

Finally, this section looks at the total number of route additions and removals by prefix size.

More Specifics

A list of route advertisements that are more specific than the original Class-based prefix mask, or more specific than the registry allocation size. There was a view at one point that routing announcements

should align with the prefix sizes that were allocated or assigned to the network. This view is no longer held within the routing community.

Possible Bogus Routes and AS Announcements

A "bogon" is a route describing an address block that is not currently assigned or allocated by a Regional Internet Registry to any network. This section of the report lists all those route objects that contain an address prefix is not currently assigned or allocated. The route descriptions also list the network AS numbers that originate these routes into BGP.

The second part of this report section lists bogon AS numbers, and the immediate upstream AS that is observed to be propagating a path for this AS.

And that's the CIDR Report.

Evaluating the CIDR Report

Has the CIDR report made a difference to the state of the interdomain routing system over all this time?

One way to evaluate this is to look at the data for more specific routes. We can look at the long-term ratio of the number of more specific routes to the total route count, and the plot of this data over the past decade is shown in Figure 9 and 10.

Specifics as % of FIB entries (552102/1063777)

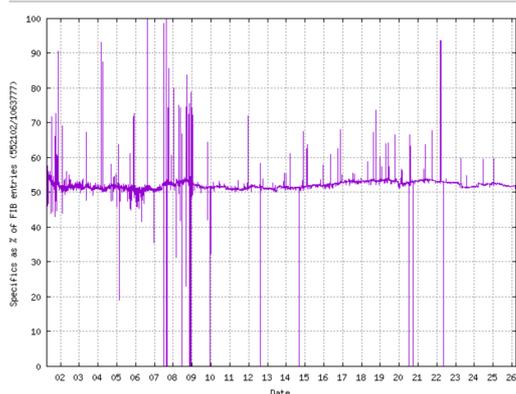


Figure 9 – Ratio of more specifics to total routes in IPv4

Specifics as % of FIB entries (140381/247168)

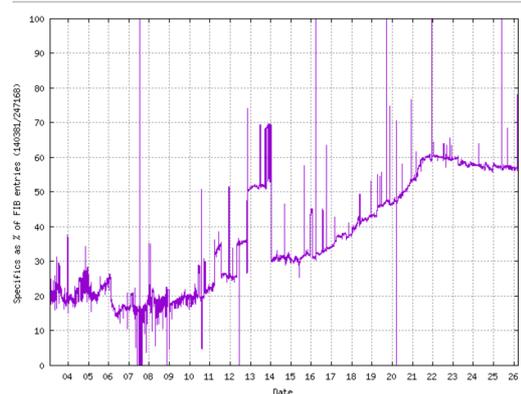


Figure 10 – Ratio of more specifics to total routes in IPv6

It is clear from these numbers that the volume of more specific routes is not changing for the better in either protocol and this leads to a conclusion that whatever traction the CIDR report may have had in the routing community more than twenty years ago, this has largely dissipated since then.

This report was the subject of a [detailed academic study](#) in 2011 by Stephen Woodrow when he was studying at the Massachusetts Institute of Technology, and a summary of this study was [presented](#) to NANOG 53 in September 2011. His conclusions were that while the CIDR report appeared to be effective in its early days, this has declined over time, and it has limited relevance and traction in the routing community in 2011 (and arguably the report has far less traction and relevance some 15 years later!)

The simple observation is that we've managed to come to terms with the Internet's large routing table, and while a smaller routing table could generate some efficiencies in routing and forwarding, the collective will to motivate each individual network to closely manage the routes that they advertise into routing table to significantly reduce the number of more specific routes is simply not present.

Part of the larger routing scaling issue was attempting to perform destination address lookups in routing hardware within the elapsed time of a processing each packet. The combination of higher line speeds and larger lookup tables exacerbated this problem. While faster hardware in routers, particularly relating to high-speed content-addressable memory, a more recent response by the vendors of large high-speed routers is to turn to performing their own form of in-router aggregation by using a technique termed "FIB compression". This is a form of proxy aggregation of route objects, where adjacent address prefixes that share a common forwarding state in the route can be aggregated into a single address prefix within the router's internal forwarding tables.

There is a further factor at play here which I would term as "the death of transit" in the Internet. These days the overwhelming volume of content and service is extensively replicated across the network using data centres located close to populations of end users. What this means is that the majority of traffic volume associated with user-triggered Internet transactions do not involve packets being passed through long haul transit routes, and the routing that manages these long haul lines is nowhere near as critical as it was a couple of decades ago. Instead, the Internet is largely being used as a collection of single hop last-mile edge networks, and the "glue" that defines a common Internet environment does not lie in a common address system or a common routing environment, but in a common name system.

So where does that leave the CIDR Report? I'd call it a largely historic artefact with little in terms of important direct relevance to the business of operating today's Internet. We've just moved on from inter-domain routing in the Internet!

Disclaimer

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

Author

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

www.potaroo.net