Counting 6to4 Relay Routers

David Malone,
Hamilton Institute
NUI Maynooth, Ireland
david.malone@nuim.ie

ABSTRACT

6to4 is a mechanism for providing IPv6 addresses and connectivity where native IPv6 is not available. In 6to4, the links between the IPv4 and IPv6 Internets are called *relay routers*. These may be advertised publicly or privately. The number of 6to4 relay routers has been the subject of debate, as additional routers increase the scalability and efficency of the 6to4 system. Counting *public* relay routers is easy using the global routing table. This paper outlines a technique that can count *private* relay routers and reports results of applying this method. Our results indicate that there are a significant number of private relays in operation in comparison to the number of public relays. This number has increased over the last two years. The results also indicate that using distributed traceroute facilities to measure the multiplicity of an anycast deployment requires large numbers of nodes to be accurate.

1. INTRODUCTION

Significant effort has been made to devise ways for people to use IPv6 in the absence of a complete IPv6 infrastructure. 6to4 is way of routing IPv6 packets over the IPv4 Internet, specified in [4]. We will just give a flavour of it here. Like tunnelling, sites using 6to4 have a router responsible for decapsulating and encapsulating packets. However, 6to4 embeds the public IPv4 address of this 6to4 router within every IPv6 address of the site, which is used for tunnelling.

The IPv6 Internet and the IPv4 Internet are joined by *relay routers*. A relay router that routes packets from the IPv6 to the IPv4 Internet advertises the 6to4 prefix, 2002::/16, into the IPv6 routing table (either locally or globally). To get an encapsulated packet from the IPv4 Internet, you send the packet to a special anycast address, 192.88.99.1[8], which may be advertised in the local or global routing table. This address can be thought of as representing the IPv6 Internet in the IPv4 network[1].

6to4 has so far proven a relatively successful IPv6 transition mechanism and there is evidence of a large number of 6to4 capable clients [14]. It should be clear that relay routers are essential to the operation of 6to4, similar to the way that the DNS root servers are essential to the operation of DNS. The more relay routers that are available, the smoother 6to4's operation will be. Consequently, the number of relay routers is something that impacts on the effectiveness of 6to4.

[1]Before this anycast address was allocated, you had to know the address of a relay router [15].

IPv4 BGP	IPv6 BGP	Name
AS559	AS559	SWITCH
	AS786	JANET
AS1741		FUNET
AS3246		SONGNETWORKS
AS8379		CYBERNET-AG
AS9033		ECIX-AS
	AS9264	ASNET
AS12859	AS12859	NL-BIT
	AS17715	CHTTL-TW
AS17832		SIXNGIX-AS-KR
	AS24895	FUBAR
AS30155		KLU

Table 1: ASs advertising 192.88.99.0/24 or 2002::/16 in v4/v6 BGP.

2. WAYS TO COUNT RELAY ROUTERS

There are a small number of 6to4 relay routers that are obvious because they advertise 192.88.99.0/24 in the IPv4 BGP routing tables. A block containing 192.88.99.1 is advertised to prevent the route being filtered out by routers that ignore small netblocks[2]. Perhaps the best known of the publicly-advertised relay routers is the one at SWITCH, the Swiss Education and Research Network, but on any day there are a number of networks offering public 6to4 relays. There are several databases collecting historical global routing information [12, 5, 6]. One snapshot from the Route Views project found that the autonomous systems (ASs) listed in the first column of Table 1 advertising 192.88.99.0/24.

However, this is not the whole story. To begin with, the list of visible relay routers in IPv6 BGP may not be the same as the list in IPv4 BGP. For example, by looking at several IPv6 looking-glasses we found the ASs listed in the second column of Table 1 advertising routes to the 6to4 prefix, 2002::/16. While the IPv6 list overlaps with the IPv4 list, it clearly isn't the same, even though they were noted contemporarily. The discrepancies may arise because of differences between IPv4 and IPv6 policy within organisations.

More interestingly, some networks may choose to provide a 6to4 relay that is only available internally, by advertising the 192.88.99.0/24 and/or 2002::/16 within their own AS (or to selected BGP peers). In such cases it is unlikely that these routes will be visible to a project like Route Views. However, if 6to4 becomes a popular method for the connec-

[2]One ISP accidently advertised 192.88.99.0/25 in the global BGP table late in 2003, drawing in many people's 6to4 traffic. The longer prefix was used internally to attract traffic.

tion of IPv6 end sites, a significant number of 6to4 users could be supported by such private relays.

So, how can we estimate the number of 6to4 relays? The ideal way would be to traceroute to 192.88.99.1 from every point in the Internet and see where those traceroutes lead. Similarly, tracerouting from points in the IPv6 Internet to some address in the 2002::/16 range would reveal the IPv6 side of 6to4 routers. This could be undertaken using a distributed facility like PlanetLab[2] or AMP[1]. Early in 2004, Matthew Luckie conducted a survey by tracerouting from the AMP nodes. This survey found around 5 relay routers, though most nodes used the well-known relay in SWITCH. Interestingly, not all the relays found in Luckie's survey were publicly advertised. Later in 2004, we conducted similar experiments using other distributed traceroute services. The Scriptroute[16] server (using PlanetLab nodes) located 7 relay routers when routing loops and duplicates were accounted for. Here most of the nodes used a relay router in SWITCH, Funet or Abilene. Similarly, the traceroute mesh server at WAND[9] found 5 relay routers, a large number being served by the SIXNGIX router, reflecting the fact that much of the mesh is in the APNIC region. To give an idea of the number of source points, AMP has about 150 nodes, Scriptroute about 250 and WAND about 60; though not all of these may be active at a particular time.

However, tracerouting from a large number of points in the Internet is not the only way to find relays. A practical way presents itself that does not require any special facilities.

3. HOW TO PERFORM A COUNT

Consider what happens if we traceroute with an IPv6 source address that is in the 6to4 range. As packets (ICMP TTL exceeded) are sent from each hop, these packets will make their way to the nearest relay router advertising 2002::/16. This router will encapsulate the packet and send it to the appropriate IPv4 address. Figure 1 illustrates this. By collecting these IPv4 packets at their destination and examining the source address used for encapsulation, we will get the addresses of the 6to4 relays serving nodes along the path that we are tracerouting.

Targets for such a traceroute are easy to find. In the IPv6 world a network provider is generally represented by a single prefix in the global IPv6 BGP tables. This routing table is still relatively compact, containing less than 1000 prefixes. By tracerouting to the first address in each range, it seems likely that a packet will make its way into the organisation's network and the ICMP reply will make its way via the nearest relay to that organisation. Performing the count described is relatively straight forward using tcpdump, traceroute and a dump of the IPv6 BGP table. The process can be completed in a few hours without stressing a modest DSL connection.

In practice, most of the returned packets are accounted for by a rather small number of encapsulating IPv4 addresses. Among these is 192.88.99.1, which may account for a number of relays. As an anycast address, 192.88.99.1 should usually not appear as a source address, however for reasons related to both operations and software, it does.

Trying to resolve those relays replying using 192.88.99.1 is important, as this may account for a significant number of relays. This can also be done with traceroute. Consider tracerouting to a 6to4 address: the last hop before the decapsulating router will be the relay router. Just as IPv4 provides loose source routing, IPv6 provides a way to traceroute via particular intermediate nodes (using routing headers). So, tracerouting to a 6to4 address *via* nodes whose relay router replied using the anycast address may reveal the IPv6 addresses associated with that relay router. Figure 2 shows an example of this.

A single relay router may have many IPv4 and IPv6 addresses. We have to consider the possibility that multiple addresses identified in the count actually belong to a single relay router. Manual inspection suggests that IPv4 addresses (other than 192.88.99.1) represented distinct relays. This is probably due to source address selection being applied consistently when encapsulation takes place for the fixed IPv4 destination used to perform the count.

For the relay routers that replied using the anycast IPv4 address, we determined their IPv6 addresses using the technique described above. This list contained groups of addresses that obviously belonged to the same router. This is due to the traceroute replies being generated by a particular interface, usually the one on which the expiring packet arrived. To account for these duplicates, only the first 32 bits of the IPv6 address were considered, and then these were checked in the whois database to eliminate duplicates (e.g. relays that use both a 6bone and production prefix).

It is worth summarising what is required to identify a relay router using these techniques. First, we must have a node that the relay router serves that also responds to one of our traceroutes. Thus the node must be along a path we are tracing and must generate ICMP Time Exceded or Port Unreachable messages. If the router encapsulates using an address other than the IPv4 anycast address we are done. Otherwise we need to be able to use the IPv6 routing header to traceroute via the node served by that relay router and we need the relay router to generate ICMP Time Exceded messages.

4. RESULTS OF THE COUNT

The count was performed in July 2003, January 2004, December 2004 and June 2005. Each count produced a number of IPv4 addresses (26, 26, 39 and 43 respectively) including 192.88.99.1. Resolving the anycast address produced a number of IPv6 /32s (12, 18, 20 and 15). We then manually accounted for duplicates and the RIPE IXP block to get a total (37, 44, 56, 57). Some relays were systematically missed because they were within two hops of the node performing the count and traceroute had been run with options to skip the first two hops, requiring a correction (1, 2, 2, 1). The breakdown in Table 2 was produced by assigning relay routers to countries using whois, traceroute and DNS. These databases are known not to be completely reliable for this purpose. Regardless, we get an indication of the geographical distribution of relays.

While it is clear that new relay routers have appeared over the course of these measurements, other relay routers do not appear in all surveys. In a small number of cases (2 or 3) it seems a relay router may have been missed by the survey. In other cases it seems more likely that the relay router has discontinued service. Overall, there has been a significant increase in the number of relays since July 2003. It is also interesting that the number of relays using 6bone addresses has decreased from 4 to 1.

Note that we can get a crude estimate of the domain of

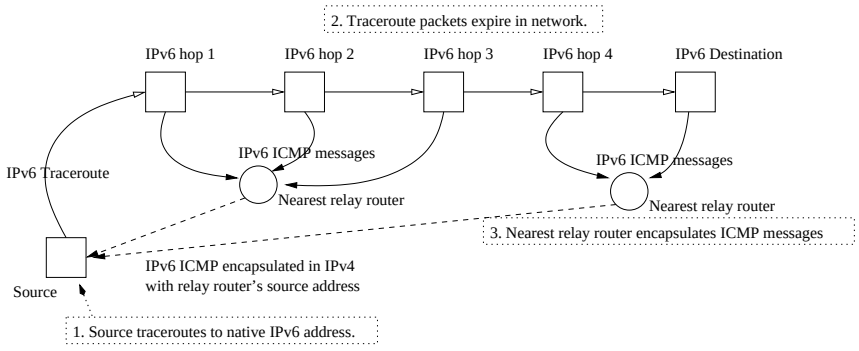

Figure 1: Identification of relay routers using IPv4 encapsulation address.

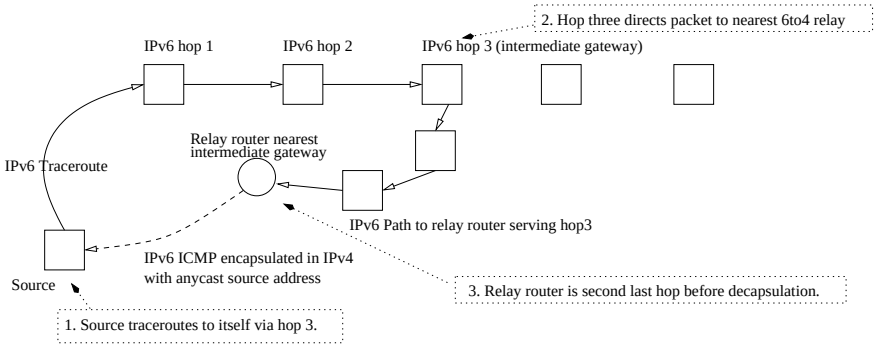

Figure 2: Identification of relay routers using traceroute and the routing header.

	AU	BE	BR	CA	CH	CN	CZ	DE	EE	ES	EU	FI	FR	GR	HU	IE	IT	JP	KR	LT	LU	MX	NL	NO	PL	PT	RU	SE	SK	TH	TW	UK	US	Tot
Jul'03	1	0	0	0	2	0	1	4	1	0	0	2	0	0	0	3	0	2	3	2	0	0	3	1	0	1	0	2	0	1	1	4	4	38
Jan'04	1	0	0	1	2	0	0	9	1	1	2	2	0	1	0	3	1	1	3	2	1	0	1	1	0	2	0	2	0	2	1	4	2	46
Dec'04	1	0	2	0	3	0	0	7	1	1	0	2	1	1	1	3	3	1	2	2	1	1	3	1	3	2	0	2	1	3	2	2	6	58
Jun'05	3	1	1	1	3	1	0	7	1	2	0	1	0	0	1	3	2	1	2	2	0	1	2	0	2	2	1	3	1	3	2	3	7	59

Table 2: Breakdown by country code.

attraction of particular routers in the IPv6 network. By examining the source IPv6 addresses of packets attributed to a particular router, we can give a lower bound on the number of IPv6 addresses, /32 networks and /48 networks served by that relay.

Table 3 shows a breakdown of the top relays. We show the country code for the relay and, for public relays, the group running it. Relays identified by their IPv6 address are shown in *italics*. Since the anycast address accounts for a large proportion of the relays, we include an aggregate of all relays responding with the anycast address for comparison. The entry for RIPE IXPs is not actually a single relay, but several relays that all use addresses in the RIPE IXP range.

Source	Packets	IPv6 addrs	/48s	/32s
Anycast	1047	407	160	103
Cisco, US	1371	118	70	52
KDD, JP	660	130	48	39
SWITCH, CH	434	121	52	29
UK	161	49	14	12
LT	27	15	13	12
RIPE IXPs	67	30	15	10
TW	52	19	9	8
EE	34	19	12	7
LT	7	7	7	7
DE	17	11	7	5
...				

Table 3: Breakdown by packet/address/network.

5. CONSIDERATIONS AND CONCLUSIONS

This counting technique has found more relay routers than the obvious technique of tracerouting from many points in the IPv4 Internet. It seems that in addition to the well-known publicly-advertised 6to4 relays, there are a number of private relays in operation. This is good news for people using 6to4, as they are more likely to get a router close to them and are more likely to have other 6to4 routers to automatically fall back to if the closest 6to4 router fails. However, it is bad news for those estimating the size of anycast populations using traceroute servers.

The list of countries found to have relays seems to be

biased towards Europe. At least part of this is likely to be systematic as the count was performed from a European IPv6 network. It is also possible that the deployment of native IPv6 is further ahead in other parts of the world, where 6to4 relays would be less common. Attempts to find relays using traceroute servers demonstrated similar bias: Scriptroute toward Abilene and WAND toward the APNIC region.

This work provides some insight into the domain of attraction in the IPv6 Internet for these relays. To get more accurate figures for the number of /48s and IPv6 addresses served, a larger number of IPv6 targets would be required. Unfortunately, the survey provides no good way to estimate the encapsulation load on each relay (though this might be estimated by observing the IPv4 ID field, or via similar techniques [13]). We do see a large number of packets being returned through a small number of relays (KDD Labs, Cisco and SWITCH).

Unlike using traceroute from a number of points in the IPv4 Internet, this survey gives no clues as to how the relays attract packets in the IPv4 network. This is one of the biggest weaknesses of this count: it is possible (though unlikely) that the relays that take packets from the IPv6 Internet to the IPv4 Internet are unrelated to the relays that send packets in the opposite direction. Table 1 suggests that there will be some overlap. The relays discovered that encapsulate using 198.88.99.1 almost certainly offer service in both directions, as there would be no reason to configure this address otherwise. Some interesting asymmetry has been noticed, where publicly advertised relays sometimes see much more traffic going in a particular direction. However, the level and direction of asymmetry seem to change over time.

There is some scope for this technique to produce incorrect results. We targeted networks with native IPv6 addressing, however it is possible that we stumbled across a single node with a native address that also had 6to4 configured. This configuration looks close enough to a network with a private 6to4 relay that we would identify it as such. It is also possible that some ICMP messages or routing headers required by our technique are filtered by firewalls within the IPv6 network. The relatively high level of consistence across the surveys suggests the methods have some robustness.

There may be some inaccuracy in the identification of particular relays, or of the network they reside in. For example, suppose a relay in one network generates an ICMP on an ingress interface and the address allocation for the link has been assigned by a neighbouring network. In this case we may attribute the relay to the wrong network, or even count it twice. Examination of the data suggests that there may be a small number of occurrences of this.

As observed above, a more comprehensive list of addresses to traceroute to might help find more relays and give better indications of the size of the populations served by each relay. The IPv6 Skitter[10] project has already collected a list of IPv6 addresses based on the 6bone database. There are plans to walk the reverse DNS tree to look for IPv6 addresses. Other possible sources of addresses included lists of IPv6 web servers that have been trawled, other IPv6 topology measurement research [7, 3] or routing projects such as Ghost Route Hunter[11].

The survey is relatively automated, the manual work involves looking for likely duplicates and querying various whois databases. A fully automated survey, recording historical information and presenting up-to-date information on the web, might provide insight into the stability of routing within the 6to4 system and would provide more certain information regarding trends in the number of 6to4 routers available.

The techniques developed to identify the relay routers might be applied in other situations. The technique used to find the IPv4 address of the relays depends on the relay encapsulating the packet (and so sending an identifier). The technique used to find the IPv6 addresses depends on being able to traceroute *via* given points in the Internet. While it seems possible to traceroute via points in the IPv6 Internet, loose source routing is often blocked in the IPv4 Internet. Nonetheless, it might be applied to count DNS root servers, where anycast deployment is common.

6. ACKNOWLEDGEMENTS

HEAnet provided snapshots of their IPv6 BGP table. Alexander Gall, Matthew Luckie and Pekka Savola provided valuable advice. This work has been supported by Science Foundation Ireland under the National Development Plan.

7. REFERENCES

[1] Nlanr's active measurement project (AMP). `http://watt.nlanr.net/`.

[2] Planetlab. `http://www.planet-lab.org/`.

[3] *Identifying IPv6 Network Problems in the Dual-Stack World*, ACM SIGCOMM Workshop on Network Troubleshooting, August 2004.

[4] B. Carpenter and K. Moore. Connection of IPv6 domains via IPv4 clouds. RFC 3056, February 2001.

[5] A. N. T. Center. University of oregon route views project. `http://www.routeviews.org/`.

[6] P. Gloor. The netlantis project. `http://www.netlantis.org/`.

[7] B. Huffaker et al. Visualizing ipv6 as-level internet topology. `http://www.caida.org/analysis/topology/as_core_network/ipv6.xml`.

[8] C. Huitema. An anycast prefix for 6to4 relay routers. RFC 3068, June 2001.

[9] P. Lorier. Traceroute mesh server. `http://wand.cs.waikato.ac.nz/~perry/tr/tr.php`.

[10] M. Luckie. IPv6 skitter. `http://www.caida.org/~mjl/`.

[11] J. Massar. Ghost route hunter. `http://www.sixxs.net/tools/grh/`.

[12] R. NCC. Routing information service (RIS). `http://www.ripe.net/projects/ris/`.

[13] S. Sanfilippo. Packet throughput guessing using the ID field of the IP protocol. `http://www.kyuzz.org/antirez/papers/ipid.html`.

[14] P. Savola. Observations of IPv6 traffic on a 6to4 relay. *ACM SIGCOMM CCR*, 35, January 2005.

[15] N. Sayer. Public 6to4 relay routers. `http://www.kfu.com/~nsayer/6to4/`.

[16] N. Spring, D. Wetherall, and T. Anderson. Scriptroute reverse path tree tool. `http://www.scriptroute.org/`.