

IS CURRENT 6TO4 RELAY DEPLOYMENT ADEQUATE?

Mónica Domingues, Carlos Friaças
Foundation for National Scientific Computing (FCCN)
Lisbon, Portugal
{ monica.domingues, cfriacas}@fccn.pt

David Malone
Hamilton Institute
Maynooth, Ireland
dwmalone@maths.tcd.ie

Abstract— *6to4 is a mechanism for providing IPv6 connectivity where native IPv6 is still unavailable. 6to4 is based on relay routers deployment. The anycast technique is used to address relays, and each relay's advertisement may be seen globally or only within a limited scope. The number of available relays is important to this mechanism's robustness, as well as their geographic location. Due to the details of the mechanism, limited scope relays are not easy to find, despite the fact their location and existence is key to evaluate suitability of current deployment.*

Keywords: *IPv6, 6to4, transition, BGP, routing.*

I. INTRODUCTION

This paper aims to discuss the state of current 6to4 relay's deployment.

The usage of this transition mechanism (like any other) only makes sense in the specific contexts. For example, where native IPv4 and IPv6 already coexist, the usage of 6to4 relays should decrease.

This paper goes through several sources and mechanisms to find 6to4 relays, with the objective of collecting knowledge to allow an evaluation about the state of deployment. The main motivation for this, about such an operational issue, is to draw attention on the negative effect that an insufficient 6to4 relay deployment may have on IPv6 in the medium/long run.

II. 6TO4 BASICS

6to4 is a mechanism defined in RFC 3056 [1] that uses a block of addresses (2002::/16) to interconnect nodes to IPv6 backbones over IPv4 networks, in a simple and seamless way for end users. It also uses the IPv4 prefix 192.88.99.0/24 [2] in order to guarantee connectivity with the IPv4 world.

Much like point-to-point tunnelling, sites using 6to4 have a router responsible for encapsulating and decapsulating packets. However, 6to4 embeds the public IPv4 address of this 6to4 router within every IPv6 address of the site. For example, a 6to4 router with IPv4 address 192.0.2.4 by definition serves the site 2002:c000:0204::/48. Thus, when a 6to4-aware node gets an IPv6 packet destined for any 6to4 address, it immediately knows the IPv4 address the packet should be tunnelled to by examining the embedded IPv4 address.

Naturally it would be better if the search for a relay could be done automatically, without having to consult a list and

issue a manual configuration. Consequently, the anycast model is used to allow the existence and automatic usage of several relays, geographically spread across the globe.

III. USING WHOIS SERVERS

Whois servers are a possible source of routing data. Usually routing information is included in whois databases to help networks design their routing policies and diagnose problems. In our case, we are specifically interested in people who list the route one has to announce, in order to advertise a public 6to4 relay. Some of the Regional Internet Registries run this kind of registry service for their communities. However, scanning for 192.88.99.0/24, we only see information on the RIPE database. Thirteen records were returned, spanning across ten countries (Table I).

TABLE I. DOMAINS ANNOUNCING 6TO4 ROUTE IN THE RIPEDB

| AS | Description | Country |
|-------|-------------|-------------|
| 559 | Switch | Switzerland |
| 1741 | Funet | Finland |
| 1930 | Fccn | Portugal |
| 2847 | Litnet | Lithuania |
| 3327 | Linxtelecom | Estonia |
| 3344 | Kewlio | U.K. |
| 9033 | Ecix | Germany |
| 12779 | Itgate | Italy |
| 12816 | Lrz | Germany |
| 12859 | Bit | Netherlands |
| 16150 | P80 | Sweden |
| 20640 | Titan | Germany |
| 29259 | labg | Germany |

Having found no data on the other four RIR whois databases, scanning the RADB [3] seemed like an obvious choice. RADB is the Routing Assets Database, and it is managed by Merit Network. This database incorporates data from other databases, including RIPE's. Consequently, we found again the same thirteen records, plus five more, from other sources, depicted in Table II.

All the records found don't really imply that the 6to4 mechanism is up and running in a correct fashion. These records are merely administrative clues to find where potential relays may be located. There is also another issue – completeness – other relays may be deployed, without any public administrative information available.

TABLE II. ADDITIONAL RECORDS FOUND IN THE RADB

| AS | Description | Country | Source |
|-------|-------------|-----------|--------|
| 7570 | Aarnet | Australia | Radb |
| 14361 | Hopone | US | Savvis |
| 17832 | Sixngix | Korea | Radb |
| 27552 | Towardex | US | Altdb |
| 29748 | Carphatia | US | Radb |

IV. SCANNING LOOKING GLASSES GLOBALLY

It is hard to locate more 6to4 relays, or even fully confirm the relays identified in the previous section. While [4] provides an interesting methodology, in this section our main line of investigation is using the resources of traceroute.org [5].

Using looking glasses to traceroute from twenty two European countries to 192.88.99.1, the output leads to several different autonomous systems. From the thirteen records shown in Table I, only seven autonomous systems were seen as the last hop of the traceroutes. Thankfully, none of the entries in Table II matches with the destination ASes found because those networks are in a different continent from their European sources!

Table III was constructed using looking glasses from other continents.

ASes containing 6to4 relays which are installed in the same continent are in bold.

The first three regions' data in Table III, show a worrying trend in 6to4 relay deployment: no relay placed in the same continent (or sub-continent) was found, having direct impact on latency. Of these three, the best case is Latin America, where we saw a minimal average roundtrip time of 40ms.

With this data, it's easy to see that in Africa and Oceania, some relays are urgently needed, and a couple of new relays in Latin America would also be advantageous. The majority of sources from where we scanned are closer to 289ms than from the minimal average value.

In the Asian and North American cases, the biggest problem is un-optimized routing. 6to4 relays exist in both continents (however Asia could benefit of some more). In North America, the existence of several relay routers should be enough to avoid transatlantic hops, which cause exaggerated RTT values.

TABLE III. LOOKING GLASSES – REGIONAL VIEWS

| Region | Views | 6to4 ASes | RTT span (ms) |
|---------------|-------|--|---------------|
| Latin America | 12 | 559, 1239, 1299 | 40-289 |
| Africa | 4 | 559, 12491, 12859, 29259 | 197-590 |
| Oceania | 15 | 559, 1239, 3356, 5050, 29748 | 188-379 |
| Asia | 19 | 559, 1239, 3356, 17832 , 20640, 29748 | 52-380 |
| North America | 45 | 559, 1239 , 3356, 5050 , 8881, 12859, 17832, 20640, 29748 | 2-300 |

Searching for the 2002::/16 route is a more difficult task. The biggest problem is finding IPv6-aware traceroute tools. From the extensive list of traceroute.org, a small list of IPv6-capable looking glasses was drawn. Due to the small nature of this list, in the future, it will be maintained at www.ipv6.eu [6], as an information resource for the community.

At this time, it is possible to see that traceroute tools allowing IPv6 functions are mostly based inside European NRENs, commercial European networks and some tier-1/tier-2 networks. These worldwide networks, which provide IPv6 service to their customers already have converted some tools for dual-stack usage. We were also able to locate IPv6-enabled traceroutes in Asia and South America. Subsequent work will look into output of these looking glasses.

V. LOOKING AT GHOST ROUTE HUNTER

The Ghost Route Hunter is a project to monitor IPv6 routing. It provides a distributed traceroute feature, which we use in this section.

In Table IV we see very low latency towards 192.88.99.1 measured from several European networks. This scenario obviously can't be generalized to all European networks, because all the networks in the table have their own 6to4 relay, or are a small number of hops from a relay.

We also note that the networks involved in the distributed traceroute feature of GRH [7] are co-managed by people with a strong interest in IPv6's operational issues. This obviously biases the results for these networks towards lower RTTs.

One other detail we draw from this data is that not every network has its own relay, but they are reaching the nearest relay possible most probably by traversing a local Internet Exchange [8].

TABLE IV. GHOST ROUTE HUNTER RTT TO 192.88.99.1

| AS | Description | Country Code | Average RTT |
|-------|-------------|--------------|-------------|
| 12871 | Sixxs noc | NL | 3.7 |
| 4589 | Easynet | BE | 26.3 |
| 2611 | Belnet | BE | 19.4 |
| 8767 | Mnet | DE | 1.3 |
| 16086 | Finnet | FI | 0.8 |
| 3344 | Kewlio | UK | 10.8 |
| 1213 | Heanet | IE | 1.5 |
| 12779 | ITGate | IT | 1.9 |
| 12859 | Bit | NL | 0.23 |
| 25137 | Nfsi | PT | 1.9 |
| 16150 | Port80 | SE | 1.3 |

VI. ROUTE SERVERS' VIEW

Route servers allow people to examine the routing table at various points in the Internet.

Traceroute.org was the chosen source to investigate route servers. Route servers can be very useful in debugging inter-domain anomalies and also to help in the fine-tuning of routes from an end-to-end perspective.

Unlike a traceroute mechanism, such as that used in Sections IV and V, information available at route servers is

broader. Using a route server, we can see how many paths there are to a single destination. In the case of the anycast prefix, one interesting thing to look at is the last AS in each path.

Table V aggregates information collected from several route servers, spread across the world, and most important in networks that carry transit packets.

The set of 6to4 relay routers observed as the best path at these route servers is likely to be a small subset of the total existing relays. This is because each router can only select one best route and the set of route servers is smaller than the set of traceroute-based looking glasses. Moreover, we see that from route servers in South Africa (SAIX) and from Australia (Optus) showing the best route to 192.88.99.0/24 ending at an AS in North America and Europe respectively. This is rather worrying.

The previous analysis could be made more complete if every internet exchange had a public route server. Receiving the routes directly from each Internet Exchange member could reveal how many members of the local community are hosting a 6to4 relay inside their network.

TABLE V. ROUTE SERVERS DATA ON 193.88.99.0/24

| AS | Description | # Paths | Best Path AS |
|------|--------------------|---------|--------------|
| 3257 | Tiscali | 1 | 33437 |
| 3549 | Global Crossing | 3 | 559 |
| 3561 | Savvis | 2 | 33437 |
| 5511 | OpenTransit | 5 | 1239 |
| 5713 | SAIX | 4 | 1239 |
| 7018 | AT&T | 19 | 1239 |
| 8220 | Colt | 34 | 1239 |
| 3582 | Oregon-IX | 53 | 1239 |
| 4323 | Time Warner | 8 | 33437 |
| 5388 | Energis | 1 | 1239 |
| 6730 | Sunrise | 2 | 559 |
| 6667 | Eunet Finland | 2 | 8767 |
| 7474 | Optus | 4 | 559 |
| 6939 | Hurricane Electric | 5 | 12859 |
| 6746 | UPC-Astral | 1 | 12816 |
| 3303 | Swisscom | 1 | 559 |

VII. PERFORMANCE ISSUES

From previous sections, we have seen that latency to a 6to4 relay can be quite high. Latency is important to 6to4 clients as all packets from a 6to4 client to the IPv6 native network must travel to the 6to4 relay. Thus latency directly impacts the performance that users of 6to4 see.

Note, the return path from a 6to4 relay may be less important, as packets making their way back from the IPv6 native network may use a completely different relay.

Note that all packets sent from a 6to4 client to the native IPv6 network must traverse one relay router. This relay element in fact acts as a translator, in an analogous way to NAT. Is it difficult to imagine users tolerating their NAT-enabled equipment on a different continent!

In contrast to the NAT-PT mechanism [9], the work done by the 6to4 relay is a lot less heavy. Importantly, a 6to4 relay is not required to keep any state. Unlike any NAT-PT equipment,

it is feasible for a 6to4 element to be in the opposite corner of the world. The main negative impact will be in the form of latency.

Another aspect to keep in mind is the connection to the IPv6 Internet itself. We've been focusing on finding the closest/best relay from several points in the world, but the next step is in fact looking into the quality of native IPv6 paths, used once the relay has been reached. It is well known that some IPv6 paths are suboptimal.

If the latency to the nearest relay (in terms of routing information, not geographically) is lower, and the network (or autonomous system) routes all packets to the IPv6 world through a high-latency tunnel, the negative impact could be as bad as choosing a distant relay. Again, if inter-continental hops are seen, the impact on latency, and thus on performance is highly detrimental.

VIII. CONCLUSIONS

Following this work, a strong recommendation is made for 6to4 relay routers deployment, especially near Internet Exchange Points. At these locations a 6to4 relay will not need or want global visibility, but will ensure optimized 6to4 access for local communities.

On another level, the build up of relays in tier-1 and tier-2 transit providers' networks would be very useful. This will provide a consistent relay to communities where a local relay is temporarily unavailable or nonexistent.

One surprising detail is to see the poor match between administrative information (whois) and what is going on at an operational level. Even when being announced within a limited scope, it's interesting and at the same time disappointing, to see that the RFC 3068 prefix is not showing up in the databases. This is important because these databases are used to construct filtering rules. People should not just worry about deploying a relay, but also about its proper announcement.

Geographically speaking, in Europe and to a smaller degree in North America, the relay deployment status seems to be in a healthy state. However, additional relays in specific places would be beneficial. It should be noted that more correctly deployed relays only have a positive impact on this transition mechanism's global deployment. In other world regions, relay deployment should be strongly encouraged, in order to provide IPv6 connectivity to communities with only one or few ISPs already linked the global IPv6 network.

Possible reasons for scarce relay deployment are security and commercial issues. In fact, this mechanism's control is limited. The main mechanism for stopping specific third party networks from using it is BGP's route handling. For commercial networks there is another show-stopper: unless carefully configured, offering a relay means routing packets for third-parties, sometimes over transit connections. The best-effort approach associated with managing 6to4 relays can also be considered an unappealing characteristic for profit-driven organizations.

The main conclusion from this work is indeed that the current deployment (end 2006) is still not adequate. With MS'

Vista rollout to happen soon, there is a real danger of a negative reaction towards IPv6 from end users, if they perceive IPv6 brings any degradation to their level of service.

ACKNOWLEDGMENT

The work from FCCN presented in this paper is partially financed by POS_Conhecimento, with the aim of promoting the usage of IPv6 services in the RCTS network, and to produce guideline documents to other organizations concerning IPv6 adoption. David Malone has been supported by Science Foundation Ireland under the National Development Plan.

REFERENCES

- [1] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.
- [2] Huitema, C., "An Anycast Prefix for 6to4 Relay Routers", RFC 3068, June 2001.
- [3] Route Assets Database, Available at URL: <http://www.radb.net>, September 2006.
- [4] Malone D., "Counting 6to4 Relay Routers", January 2006, in press.
- [5] Kernen, T., Traceroute.org, Available at URL: <http://www.traceroute.org>, September 2006.
- [6] The European IPv6 Task-Force Steering Committee Portal, Available at URL: <http://www.ipv6.eu>, September 2006.
- [7] Ghost Route Hunter, Available at URL: <http://grh.sixxs.net>, September 2006.
- [8] European IX Association, Available at URL: <http://www.euro-ix.net>, September 2006.
- [9] Tsirtsis G., Srisuresh P., "Network Address Translation - Protocol Translation (NAT-PT)", RFC 2766, February 2000.
- [10] Savola P., Patel C., "Security Considerations for 6to4", RFC 3964, December 2004.
- [11] Huitema, C., Austein, R., Satapati, S., Van der Pol, R., "Evaluation of IPv6 Transition Mechanisms for Unmanaged Networks", RFC 3904, September 2004.
- [12] Savola P., "Observations of IPv6 Traffic on a 6to4 Relay", January 2005, in press.
- [13] Friaças C., Baptista M., Domingues M., Ferreira P., "6to4 versus TunnelBrokers", August 2006, in press.
- [14] Hei Y., Kubota, A., Hotta T., Yamazaki K., "Open 6to4 Relay Router Operation", January 2003, in press.
- [15] Hei Y., Yamazaki K., "Traffic Analysis and Worldwide Operation of Open 6to4 Relays for IPv6 Deployment", January 2004, in press.