Flexible Routing with Maximum Aggregation in the Internet

Pedro A. Aranda Gutiérrez

University of Paderborn, Germany
paaguti@hotmail.com

Abstract. The explosion of the Internet's routing tables has been a concern in the last years. Specially after IANA assigned the last /8 prefixes on the 3^{rd} of February, 2011, two fronts are open for the Internet community: the growth of the IPv4 routing table due to fragmentation introduced by the last assignments made by RIRs and the strategy to follow for the new IPv6 Internet. This paper analyses the behaviour of the IPv4 routing table in the Internet's Default Free Zone in 2010 and presents the evolution and the current status of the IPv6 routing table in the DFZ. These paper also presents a prototype implementation of the routing architecture based on parallel routing tables. This prototype implementation was tested in an emulated environment using Netkit. This implementation demonstrates that parallel routing tables are an easy and clean alternative to current practises in order to avoid routing configurations that intend to have effect on a scoped area of the Internet are leaked outside it. This characteristic makes parallel routing tables a good candidate for Traffic Engineering configurations in IPv6.

Keywords: Routing protocols, Network Operations, Network management, Network monitoring.

1 Introduction

The explosion of the IPv4 routing table in the Default Free Zone (DFZ) of the Internet continues to be a threat, even after the Internet Assigned Numbers Authority (IANA) handed out of the last /8 prefixes to the Regional Internet Registrys (RIRs) the 3^{rd} of February, 2011 [9]. The same concern is growing regarding the IPv6 address space. This is reflected in the strict policies the RIRs are imposing on IPv6 address allocations [4]. I share this concern and have proposed to use parallel routing tables in the Internet in order to isolate the Internet's DFZ from Traffic Engineering (TE) artifacts. This solution makes it possible to apply current practises in TE and keep maximum aggregation in the DFZ. It would be an enabler for a quicker adoption of IPv6. Adoption of IPv6 is a major concern, now that the last /8 prefixes held by IANA were handed out and some regions of the Internet (i.e. America, APAC and Europe) face IPv4 address space depletion in the near future. In this paper, I present a prototype implementation of the routing architecture using parallel routing tables based

K. Pentikousis et al. (Eds.): MONAMI 2011, LNICST 97, pp. 332–343, 2012.

on the open-source Quagga Routing Suite [20]. To check the properties and viability of the implementation, a proof-of-concept testbed using Netkit [27] has been used.

The rest of this paper is structured as follows: Section 2 analyses the latest trends in the growth of the IPv4 routing table in the light of new findings [13] and examines how the IPv6 routing table is behaving. Section 3 presents the prototype implementation for the routing architecture proposed in MONAMI-2010 and compares how it behaves with other setups that can be considered current practises. Section 4 presents related research and Section 5 presents the conclusion and future work.

2 Evolution of the Internet

The Internet is entering a transition phase it has long tried to avoid. Since the 3^{rd} of February, 2011 it is clear that the IPv4 address space is facing exhaustion and that IPv6 needs to be deployed. In this section, the evolution of IPv4 over the last 10 years and the evolution of IPv6 are presented and studied under the perspective of aggregation. Besides exhaustion, the second most important problem faced by the Internet is an explosion of the routing table size in the Default Free Zone, understanding by explosion an evolution that overwhelms the technology in terms of memory and processing capacity.

2.1 Evolution of IPv4

The routing table for the IPv4 routing protocol is continuously growing in the Internet's Default Free Zone. Figure 4(a) shows the evolution of the routing table size collected by the RIPE's Routing Repository (RIPE RR). The graph takes data from collector RRC00, situated at the RIPE-NCC's DFZ area. It shows steady growth stretching through 2010 despite the economic downturn. The outlier in 2008 is due to failures in the collecting procedure which have been documented by Cheng et al. in [5]. In [1], I proposed an algorithm to assess the fragmentation in the address space of the Internet's DFZ. This algorithm compresses routing tables by looking for disaggregated prefixes advertised by an ISP. These prefixes are then substituted by the next better aggregation (i.e. two adjacent /24 prefixes are aggregated to their common /23 super-network). The algorithm is recursive and most aggregation is obtained in the first steps. Figure 1 shows the evolution of size of the routing table in the DFZ between January, 2001 and December, 2010. Figure 2 shows the evolution of the aggregation achieved with the first three iterations of the proposed , expressed as the percentage of routes that could be eliminated from the original routing table.

It shows how, Between 2002 and 2009, this aggregation ratio grew lineally, but during 2010, it remained constant. Possible explanations for this behaviour, that have been given are:

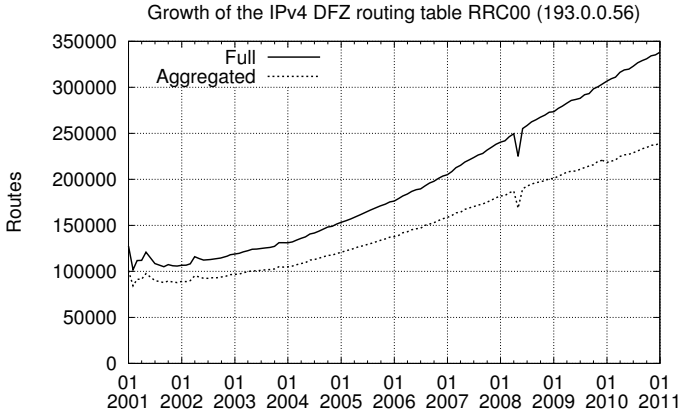

Fig. 1. Evolution of the IPv4 DFZ: Routing table size

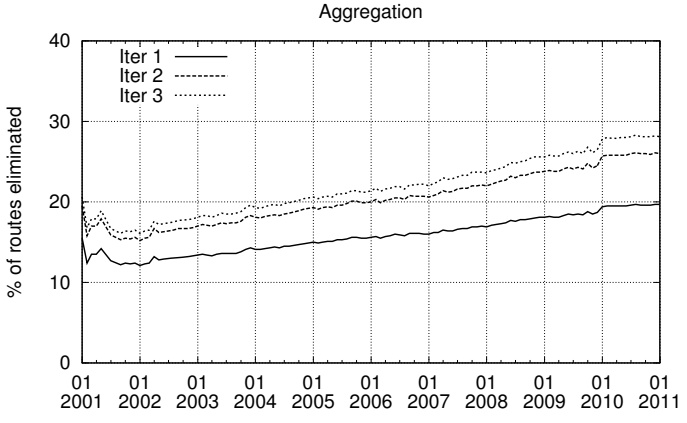

Fig. 2. Evolution of the IPv4 DFZ: Disaggregation

- *the deep economic crisis* that started around that time, *which would have slowed down the growth of the Internet*. However, Figure 1 does not suggest that this happened: during 2010, the IPv4 routing table continued to grow. Moreover, as Figure 4(a) shows, the number of leaf Autonomous Systems (ASs), i.e. ASs that advertise prefixes to the Internet, continued to grow during 2010 as in the previous years.
- *the depletion of the IPv4 routing space that has forced the RIRs to allocate smaller prefixes to ASs*. This would translate in less possibilities to fragment the address space, given that the smallest prefixes that can be advertised to the Internet are /24 [3].

A third explanation for this change of trend in the aggregation rate of the routing tables of the DFZ of the Internet could also be the transformation observed by

Labovitz et al. in [13]. In this recently published paper, the authors argue that the structure of the Internet has changed radically. Some of the ASs in the core of Internet have experienced out-bound traffic growth because they host the most popular applications, sites, etc.. The core ASs have evolved from simple traffic exchanges to traffic sources. Thus, they are no longer interested in controlling their input traffic and could be reducing the number of prefixes they advertise, thus stabilising the dis-aggregation ratio. This problem has been passed to the new consumer ASs, who are charged by the volume they consume. As of writing this paper, another move to consolidate the core of the Internet has happened with the merger of two major Internet players: Global Crossing and Level 3 [6]. It remains to be seen how this merger will affect the structure of the core of the Internet.

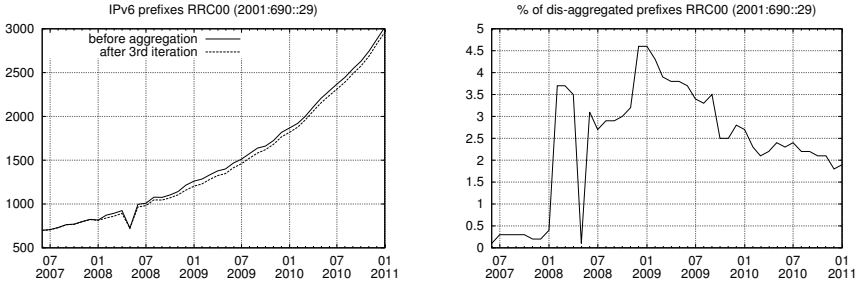

(a) Prefixes prior and after aggregation (b) Percentage of dis-aggregated prefixes

Fig. 3. Evolution of the IPv6 routing table as collected by RRC00

2.2 Evolution of IPv6

One of the main fears in the community is that the evolution of the IPv6 routing tables mimics that of the IPv4 routing table once the new protocol takes up. This is reflected in the current policy documents of the RIPE [4], where a lot of stress is put on aggregation. Figure 3 shows the evolution the number of prefixes in the Default Free Zone of the IPv6 Internet, the size of the resulting routing table after 3 iterations and the percentage of routes the algorithm was able to aggregate between 2007 and 2010. The number of routes is still quite low to draw solid conclusions. However, a very small fraction of ASs dis-aggregating their prefixes can be observed (around 50 routes or around 2% of the total routing table). Whether this level of aggregation is maintained or not depends on the number of ASs using disaggregation as part of their policies. The challenge for the IPv6 community is that the routeable address space is 48 bits[1] long or 2^{24} times greater than in the current Internet.

[1] IPv6 addresses are 128 bits long, but the least significant 64 bits have been reserved for the end user. The IPv6 equivalent to an IPv4 address is, thus, a /64 prefix. In IPv4, the smallest route-able prefix is a /24 prefix; equivalently, in IPv6, the smallest route-able prefix is a /48 prefix.

2.3 Comparative Behaviour

Figures 4 and 5(a) show that the a comparison between IPv4 and IPv6 is not possible at this point in time. IPv4 is a mature protocol, while efforts to migrate to IPv6 are starting to be seen in the community. Nonetheless, lessons learnt in IPv4 are valid for IPv6. One of the good news in the current status of IPv6 is that the majority of leaf ASs are well-behaved and only advertise one prefix to the IPv6 DFZ, as shown in Figure 5(a). This is far from happening in IPv4. As shown in Figure 4(a), IPv4 leaf ASs advertise a mean of approximately 10 prefixes per AS. This ratio has remained constant over the last years.

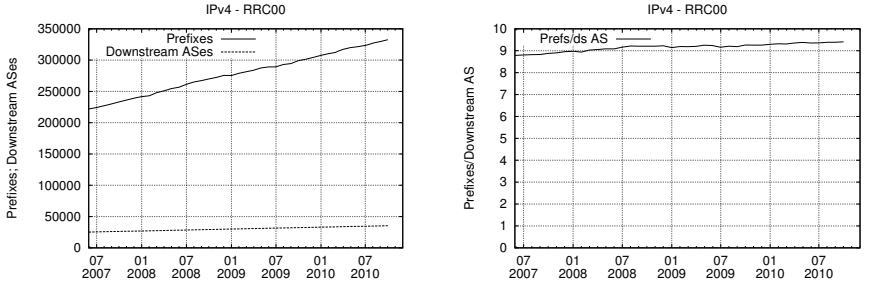

(a) Evolution of Prefixes and Leaf ASes (b) Evolution of the prefix/leaf AS ratio

Fig. 4. Prefix per leaf AS ratio in the IPv4 DFZ

3 Routing for Maximum Aggregation: A Prototype Implementation

In [1], I proposed to control aggregation in the IPv4 routing tables by making sure that only the best aggregations were present in the DFZ and that the disaggregation introduced for TE or security purposes should be kept local to the routers it was meant for. To that avail, I proposed to use parallel routing tables.

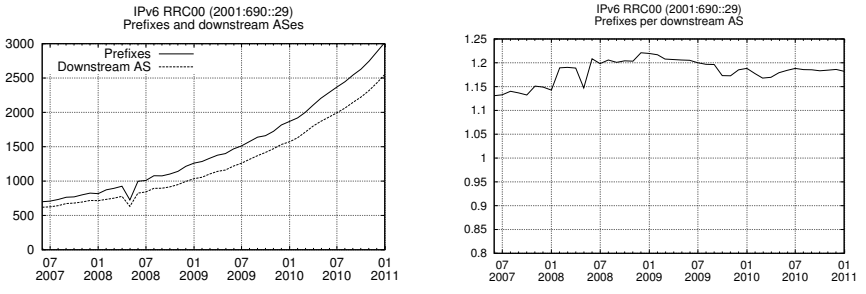

(a) Prefix per leaf AS ratio in IPv6 (b) Prefix per leaf AS ratio in IPv6

Fig. 5. Prefix per leaf AS ratio in IPv6

3.1 Prototype Implementation

Figure 6 shows the proof-of-concept implementation of a router implementing parallel routing tables.

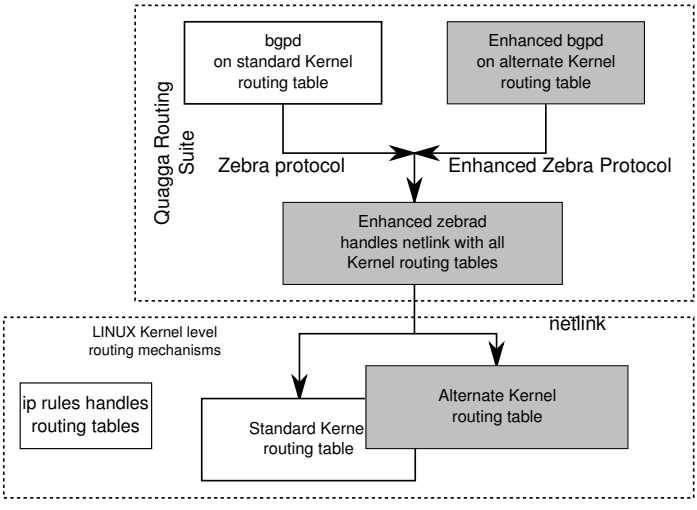

Fig. 6. Prototype implementation of the proposed routing architecture

The implementation is based on the 0.96.16 code base of the Quagga routing suite [20], an open source fork of the Zebra routing suite [10]. Both have been implemented with multiple operating systems in mind. They have a modular implementation, with a central module implementing an abstraction layer for the routing mechanisms provided the target system known as the ZServ API and different routing protocol daemons. At this point in time Quagga supports the following IPv4 and IPv6 routing protocols: RIPv1, RIPv2, RIPng, OSPFv2, OSPFv3, IS-IS, BGP-4 and BGP-4+. Additionally, external projects have implemented other protocols like LDP [24].

In order to access any routing table managed by the Linux kernel, the Zserv API was extended. An extra field carrying the kernel table identifier was introduced in the functions that manipulate the routing tables. The modification is backwards compatible: a flag indicates whether the kernel table index is included and when not, the default routing table is assumed.

In order to keep the modifications to the BGP-4 daemon to a minimum, the implementation uses two BGP-4 daemons that run in parallel. One uses the standard BGP-4 port and the standard vty port defined by Quagga and this daemon handles the main Internet routing table with the best aggregations in the Linux kernel's main routing table. The second daemon uses non-standard ports and handles the disaggregated prefixes on a separated kernel routing table. The kernel routing tables are integrated using the 'ip rules' command at system level.

3.2 Proof of Concept Testbed

The development and tests of the modified Quagga and a proof of concept were implemented in a Netkit [19, 27] environment. The topology is shown in Figure 7. It follows the general principle of a layered three-tier topology observed by Labovitz et al. in [13] in the current IPv4 Internet. The central core layer is implemented by four fully meshed ASs, AS#100 to AS#103, implemented with a single router. For the sake of simplicity, mono-router ASs are shown with their Autonomous System Number (ASN) only. The second layer is implemented AS#1000, AS#1001 and AS#1003 with a single router and AS#1002 with four border routers (r_10021 through r_10024) and a route reflector (rr_1002). The third layer is implemented again with single router ASs (AS#1010 and AS#1011) and AS_1012 with one router and two hosts. The comparison between current disaggregation practises and the proposed architecture based on parallel routing tables were implemented in AS_1002 and AS_1012 for comparison.

3.3 Traffic Balancing Techniques: A Comparison

The proof-of-concept network emulation environment was used to compare different traffic balancing techniques, that can be considered current practises. Two different scenarios were examined:

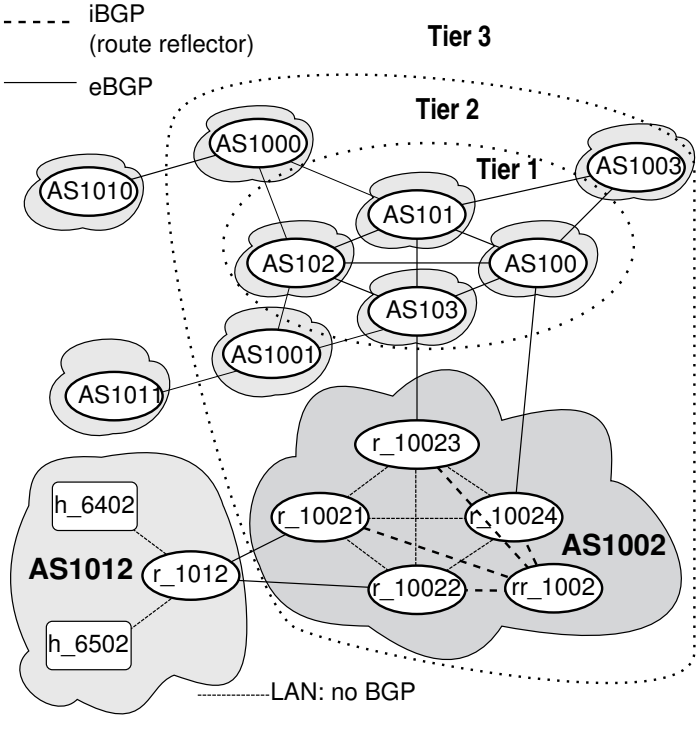

Fig. 7. Proof of concept topology

1. Stub AS with first upstream (Tier 1) AS
2. Stub AS with Tier 0 AS

The study of the stub AS case has been performed in other instances [25]. Taking into account the consolidation process in the Internet, the number of Tier1 and Tier2 ASs connected to one provider will grow. In all cases, the stub AS advertised a /20 prefix with its 16 /24 prefixes. Table 1 shows the different techniques used implement traffic balancing. The /24 prefixes were marked to use a given link as primary or secondary link. When using Multi-Exit Discriminator (MED), the upstream AS signals the priority of the whole link to the leaf AS.

Table 1. Different traffic engineering techniques used

Mark using	Primary link	Secondary link
Well-known communities	advertised with community NO_EXPORT	not advertised
Multi-Exit Discriminator	upstream AS marks complete link with "better" MED	upstream AS marks complete link with "worse" MED
AS_PATH Prepending	advertised with shorter AS_PATH	advertised with longer AS_PATH

Stub AS with Tier 1. The techniques of Table 1 were compared with the proposed architecture. The criteria used for this comparison were whether the sub-nets are advertised in the Internet's DFZ or not, whether during this process they keep the metric information for use further upstream, whether operation and maintenance procedures may result in accidental leak of prefixes, and whether traffic balancing can be implemented using the technique or not.

Table 2. Comparison between different BGP-4 control techniques

	Subnets leaked to the Internet	Subnets keep metric	Operations can result in leak	Balancing Implemented
MED	Yes	No	N/A	No
Well known communities	No	N/A	Yes	Yes
AS_PATH Prepend	Yes	Yes	N/A	Yes
Parallel Routing Tables	No	N/A	No	Yes

Table 2 shows the comparison between the different techniques. It can be argued what *quality criteria* to use in this classification. I prefer either not to advertise at all, or making sure that once a prefix is advertised, the routing information is correctly mapped to my preferences for inbound traffic. In this sense, well known communities do not provide a good solution. Regarding routing table size growth in the Default Free Zone, AS_PATH Prepending performs worse than the proposed architecture. The same applies for MED and the NO_ADVERTISE

community in case of misconfiguration. Last but not least, it also has to be remarked that the use of MED is not traffic balancing technique, bu rather a way for the upstream AS to impose traffic flows on the leaf AS.

Stub AS with Tier 0. In the case of the interaction of AS1012 with a Tier 0 provider, only *AS_PATH Prepending* can be applied in order to control the traffic coming from it. AS103 was chosen for the proof of concept. In this case, the two last lines of Table 2 hold.

4 Related Work

This paper continues work previously presented in [1]. In that paper I concentrated on the evolution of aggregation in the IPv4 routing tables until 2009. This paper continues the work with an analysis of 2010 under the light of Labovitz's observations of the evolution of the Internet. It presents a practical implementation of the routing architecture based on best aggregations that respects the address allocations made by the RIRs I proposed. By choosing this approach, the mapping between AS and prefix is respected. This is very important when debugging the Internet.

Different algorithms and approaches to compress either the Forwarding Information Base (FIB) or the Routing Information Base (RIB) have been proposed. One of the first attempts was presented by Draves et al. in [7]. FIB compression has been retaken recently by Liu et.al. in [16]. They retake the original OTRC algorithm and apply it to DFZ routing tables collected in 2009 and show that FIB compression continues to be a feasible approach to contain the look-up times in today's Internet. However, it does not attack the routing table explosion problem. Other work related to the compression of the Internet's core routing table includes the Virtual Aggregation (VA) proposal ViAggre [2]. Virtual aggregation is one of the working items of the Global Routing Working Group (GROW) in the IETF and is currently being extended to multi-AS configurations. Coupling FIB with RIB compression and extending it to the Internet has been proposed by Khare et.al. in [12]. This paper argues that FIB compression techniques can coexist with RIB compression techniques like VA and that VA can be extended beyond the AS borders. The approach presented in this paper is more natural and easier to adopt by Internet Service Providers (ISPs) since the routing tables do not loose their current look and feel. VA would require ISPs to *learn* the new mapping. Other recent attempts to modify the behaviour of Border Gateway Protocol (BGP-4) in order to make to more scalable and predictable include the proposal of imposing *next hop routing* on the Internet and getting rid of the Autonomous System Path (AS_PATH) made by Shapira et al. in [21]. This approach is even more radical than the architecture proposed in this paper. One of its merits is getting rid of AS_PATH artifacts.

Other implementations of BGP-4-based TE solutions have also been discussed by Uhlig and Bonaventure in [25] and [26]. There have been attempts at enhancing BGP-4 and limiting the topological scope of advertisements. Li et al. tried

to introduce the AS_PATHLIMIT attribute [15], meant to suppress certain advertisements after the AS_PATH attribute has reached a certain length, never passed beyond the Internet draft status. It was included in the Quagga Routing Software suite. However, the change logs for recent versions show that this attribute is no longer recognised by it.

Separating a BGP-4 into independent sessions in order to improve the isolation between the different address families running on a router was already proposed by Scudder et al. in 2003 [22]. Other closed discussions in this area have proposed to reuse the well known TCP port [28] for multiple BGP-4 sessions. Multisession BGP-4 is currently being revisited [23]. The current version of the draft proposes to use different sessions for different Address Families, but does not propose to separate the handling these Address Families by different processes. The prototype implementation includes this step. Complete separation provides protection of the BGP-4 information carried in the different sessions. However, it does not protect against the impact of an unstable BGP-4 process on the FIB.

During discussions of the Internet Architecture Board (IAB) regarding the scalability of the routing tables in the Internet's DFZ in 2006, a new change of paradigm was proposed. Some proponents argued that IP addresses are currently used both as Routing Locators (RLOCs) and as Endpoint Identifiers (EIDs) and that this duality needed to be broken. This discussion is known generically as the "Locator/ID (Loc/ID) split". Different implementation proposals have been presented. The Locator/ID Split Protocol (LISP) [8], which implied no modifications in the host protocol stack, was proposed for the last time in March, 2009 and has been abandoned. Another proposal, the Host Identity Protocol [18] has reached the Request for Comments (RFC) status and is being proposed in the context of IPv6 and the migration to IPv6. However, all Loc/ID solutions exhibit several architectural issues [17], including the fact that all solutions rely on BGP-4 to carry the information and thus exhibit the same problems of BGP-4 like the possibility of injecting bogus routes to divert traffic. Although the parallel routing tables architecture I propose doesn't solve this problem either, these routes can be detected more quickly than today: the current countermeasure for spoofed BGP-4 routes is dis-aggregation. If an attacker sends a spoofed /24 prefix, the attacked AS sends it too and leaves it to the BGP-4 route decision process to choose between the rightful and the spoofed advertisement. The rightful advertisement will "win" in some ASs, while the spoofed will be chosen in others. This makes debugging more difficult. With my proposal, the spoofed advertisement will be installed all over the Internet and thus detection mechanisms [11, 14] will deliver consistent results confirming the attack.

5 Conclusion and Further Work

This paper has revisited the growth of the IPv4 and IPv6 routing table in the Internet's Default Free Zone. It shows that current Internet trends towards consolidation of content producers [13] and depletion of the IPv4 routing space are resulting in a slow-down of the disaggregation rate of the DFZ. The current trend in IPv6 looks promising. Possibly because there is no real need for dis-aggregation

and because current policies stress on aggregation [4], the disaggregation level in the IPv6 routing table is very small. Nonetheless, the IPv6 protocol has not taken up as expected and it remains to be seen if the IPv6 routing tables in the Internet's DFZ will continue to behave like this in the future.

A prototype implementation of a routing architecture based on parallel routing tables presented in MONAMI-2010 has been presented and compared with current practises. The initial results regarding operation simplicity indicate that this architecture might help reducing the operation complexity.

Future work includes a long term observation of aggregation trends in the IPv6 Internet, once it starts to take up. It should be interesting to see how large scale adoption affects the IPv6 DFZ and what current BGP-4 practises are adopted for IPv6 operations, taking into account that multi-homing practises as known today are not favoured by current policies [4]. In this context, routing with parallel routing tables could be used to replicate most of today's TE techniques that rely on dis-aggregating prefixes locally while hiding them from the global IPv6 routing tables. It could therefore become an enabler for IPv6 take-up.

Acknowledgement. This work was made possible by the extensive BGP-4 data collections of the RIPE's Routing Repository and because NetKit is available to all of us. Without it, the implementation, test and proof of concept would have been much more difficult. Grazie tante!

References

1. Gutiérrez, P.A.A.: Revisiting the Impact of Traffic Engineering Techniques on the Internet's Routing Table. In: Pentikousis, K., Agüero, R., García-Arranz, M., Papavassiliou, S. (eds.) MONAMI 2010. LNICST, vol. 68, pp. 26–37. Springer, Heidelberg (2011)
2. Ballani, H., Francis, P., Cao, T.: ViAggre: Making Routers Last Longer! In: Seventh ACM Workshop on Hot Topics in Networks. ACM (November 2008)
3. Bush, R., Carr, B., Karrenberg, D., O'Reilly, N., Sury, O., Titley, N., Yilmaz, F., Wijte, I.: IPv4 Address Allocation and Assignment Policies for the RIPE NCC Service Region. RIPE Address Policy Working Group Document ripe-492, RIPE (February 2010)
4. Carr, B., Sury, O., Martinez, J.P., Davidson, A., Evans, R., Yilmaz, F., Wijte, I.: IPv6 Address Allocation and Assignment Policy. RIPE Address Policy Working Group Document ripe-512, RIPE (February 2011)
5. Cheng, P., Zhao, X., Zhang, B., Zhang, L.: Longitudinal study of BGP monitor session failures. SIGCOMM Comput. Commun. Rev. 40, 34–42 (2010)
6. Daneman, M.: Global Crossing sold for 3 Billion Dollars to Level 3 Communications (April 2011), `http://www.democratandchronicle.com/article/20110412/BUSINESS/104120311/0/PODCAST07/Global-Crossing-sold-3B-Level-3-Communications?odyssey=nav|head`
7. Draves, R., King, C., Venkatachary, S., Zill, B.D.: Constructing Optimal IP Routing Tables. In: Proc. IEEE INFOCOM, pp. 88–97 (1999)
8. Farinacci, D., Fuller, V., Meyer, D., Lewis, D.: Locator/ID Separation Protocol (LISP). Internet-Draft draft-farinacci-lisp-12, Internet Engineering Task Force (March 2009) (expired)

9. IANA IPv4 Address Space Registry (February 2011), http://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xml (last visit March 30, 2011)
10. Ishiguro, K.: GNU Zebra (2003), http://www.zebra.org (last visit April 09, 2011)
11. INTERSECTION (INfrastructure for heTErogeneous, Resilient, SEcure, Complex, Tightly Inter-Operating Networks) (January 2008), http://www.intersection-project.eu/ (last visit June 25, 2010)
12. Khare, V., Jen, D., Zhao, X., Liu, Y., Massey, D., Wang, L., Zhang, B., Zhang, L.: Evolution towards global routing scalability. IEEE Journal on Selected Areas in Communications 28(8), 1363–1375 (2010)
13. Labovitz, C., Iekel-Johnson, S., McPherson, D., Oberheide, J., Jahanian, F.: Internet inter-domain traffic. In: SIGCOMM 2010, pp. 75–86 (2010)
14. Lad, M., Massey, D., Pei, D., Wu, Y., Zhang, B., Zhang, L.: PHAS: A Prefix Hijack Alert System (2006)
15. Li, T., Fernando, R., Abley, J.: The AS_PATHLIMIT Path Attribute (2001), http://tools.ietf.org/html/draft-ietf-idr-as-pathlimit-03 (last visit: January 17, 2010)
16. Liu, Y., Zhao, X., Nam, K., Wang, L., Zhang, B.: Incremental Forwarding Table Aggregation. In: GLOBECOM, pp. 1–6. IEEE (2010)
17. Meyer, D., Lewis, D.: Architectural Implications of Locator/ID Separation. Internet-Draft draft-meyer-loc-id-implications-01, Internet Engineering Task Force (January 2009) (expired)
18. Moskowitz, R., Nikander, P.: Host Identity Protocol (HIP) Architecture. RFC 4423 (Informational) (May 2006)
19. Pizzonia, M., Rimondini, M.: Netkit: easy emulation of complex networks on inexpensive hardware. In: de Leon, M.P. (ed.) TRIDENTCOM, p. 7. ICST (2008)
20. Quagga Routing Suite (December 2009), http://www.quagga.net (last visit April 09, 2011)
21. Schapira, M., Zhu, Y., Rexford, J.: Putting BGP on the Right Path: A Case for Next-Hop Routing. In: Proceedings of the Ninth ACM Workshop on Hot Topics in Networks. ACM (2010)
22. Scudder, J., Appanna, C.: Multisession BGP. Internet-Draft draft-scudder-bgp-multisession-00, Internet Engineering Task Force (November 2003) (expired)
23. Scudder, J., Appanna, C., Varlashkin, I.: Multisession BGP. Internet-Draft draft-ietf-idr-bgp-multisession-06, Internet Engineering Task Force (March 2011) (work in progress)
24. Sourceforge: MPLS-Linux project (November 2009), http://sourceforge.net/apps/mediawiki/mpls-linux/index.php?title=Main_Page
25. Uhlig, S., Bonaventure, O.: Designing BGP-based outbound traffic engineering techniques for stub ASes. Comput. Commun. Rev. 34 (2004)
26. Uhlig, S., Quoitin, B.: Tweak-it: Bgp-based interdomain traffic engineering for transit ass. In: Proc. Next Gen. Internet Networks, pp. 75–82 (2005)
27. Università Roma Tre; Computer Network Laboratory. Netkit: The poor man's system to experiment computer networking (December 2009), http://wiki.netkit.org/index.php/Main_Page
28. Varlashkin, I.: Multisession BGP extensions without new TCP ports. Internet-Draft draft-varlashkin-idr-multisession-same-port-00, Internet Engineering Task Force (April 2010) (expired)