2014 IEEE 28th International Conference on Advanced Information Networking and Applications

ASN-FWD: Shrinking the IPv4 Share on the Forwarding Information Base

Marta C. C. Lacerda*, Marcos Siqueira†, Paulo R. S. L. Coelho*, Luis F. Faina*, Lásaro J. Camargos*,
Christian Esteve Rothenberg† and Rafael Pasquini*
* Faculty of Computing (FACOM/UFU), Uberlândia – MG – Brazil
Email: marta@cti.ufu.br, faina@ufu.br, {paulo, lasaro, pasquini}@facom.ufu.br
† School of Electrical and Computer Engineering (FEEC/UNICAMP), Campinas – SP – Brazil
Email: marcos.siqueira@gmail.com, chesteve@dca.fee.unicamp.br

***Abstract*—This paper presents a proposal for shrinking the number of IPv4 FIB (Forwarding Information Base) entries required on routers. Traffic forwarding under the proposed mechanism is based on the current ASNs (Autonomous System Numbers), and can be gradually adopted by ISPs. We find that, at the cost of adding 8 bytes per packet, the proposed ASN-FWD technique is capable of providing full IPv4 traffic forwarding based on ASN information, which is correspondent to 10% of the current number of IPv4 prefixes present on FIB of routers. Among its main benefits, the proposed approach alleviates the pressure on the amount of FIB shipped on routers, and paves the way for a worldwide adoption of IPv6.**

***Index Terms*—BGP, FIB, Internet, IPv4, IPv6, Routing.**

I. Introduction

Telecom equipment vendors are continuously pushed to evolve their platforms to sustain the growing Internet traffic. While optical transmission capacity is being increased very quickly, with many 100Gbps systems being deployed today, and commercial 400Gbps and 1Tbps becoming available in a few years, large and scalable IP routers also are being promoted by main vendors. These equipments commonly scale based on to multi-chassis solutions, with distributed switching fabrics and powerful route processors. The challenge of performing packet-by-packet forwarding, using LPM (Long Prefix Match) based on full route BGP (IPv4 and IPv6) routing tables (currently nearing 500K entries [1]) has been faced with flow-based forwarding mechanisms, such as Cisco Express Forwarding (CEF) [2].

Most of backbone routers are already prepared to work with dual-stack. During the transition period to IPv6, which may be very long, routers will need to manage to work with duplicated mechanisms such as routing tables and forwarding mechanisms. Most of the current high capacity routers support one million IPv4 entries shared with IPv6 on common FIB (Forwarding Information Base) memory structures. For example, such a high capacity router can be configured to store 512K IPv4 entries and 256K IPv6 entries. Therefore, based on the growth of routing tables [3], it can be predicted that in short to mid term, there will be a need of equipment upgrades. From a technical perspective, these upgrades may drive the affected backbone routers into various engineering limits, and from business perspective, networks will turn less cost-effective given the rise of the per packet (commodity) forwarding price.

Besides the migration to IPv6 that raises by itself several challenges in the Internet, the current IPv4 inter-domain routing has been facing problems due to the number of entries (IP Prefixes) required in the routing tables [3]. Such problems are caused by a number of factors, including but not limited to mobility, multi-homing, and the use of IPv4 Provider Independent (PI) pool of addresses. These issues have been investigated over the last decade by different (clean-slate and evolutionary) proposals [4], [5], [6], [7]. A common shortcoming of these proposals has been deployability, since most proposed solutions in the past require a mix of changes in the routing architecture, addressing schemes, new infrastructure devices, and so on, reducing the chances of operational adoption.

This paper proposes ASN-FWD (Autonomous System Number-based ForWarDing) which aims to allow DFZ (Default-free zone) routers to populate their FIB with 32-bit ASN entries, instead of IP Prefixes. The main advantage of ASN-FWD is the reduction in the current number of IPv4 entries, while avoiding the explosion of IPv6 entries in the near future. ASN-FWD is capable of providing full (worldwide) IPv4 traffic forwarding based on a number of entries on FIB which represent approximately 10% of the current number. While approaches with ASN routing have been proposed before [8], our proposal is unique, since we show a path towards smooth migration, providing retro-compatibility with the current mechanism based on IP Prefixes. ASN-FWD reduces technical and business challenges, creating a scenario were one million entries on FIB are still sufficient for the global scale Internet routing on the upcoming years, reducing CAPEX and OPEX of carriers not only during the IPv4 to IPv6 migration, but also in the long term, since ASN-FWD can also be used for IPv6 traffic forwarding.

The proposed ASN-FWD mechanism can be implemented by so-called *adaptation-boxes* located at source and destination ISPs. The mechanism works by replacing the 32-bit destination IP address with the destination AS number (ASN), which is extracted from BGP AS-Path information in the BGP database. In order to allow the destination ISP adaptation-box to forward the packet to the correct destination host, the

1550-445X/14 $31.00 © 2014 IEEE
DOI 10.1109/AINA.2014.161

original destination IP address is maintained in the packet by adding an 8-bytes-long optional header. Therefore, the DFZ routers only need to know about the /32 ASN entry, instead of all the IP Prefixes owned by the given ASN. This paper details the proposed mechanism and presents an analysis about its adoption possibilities in the current Internet. The overall discussion is based on the predicted growth of the FIB reported at [9], describing the time window existent for the full deployment of ASN-FWD and the impacts of using it. In addition, experimental results of the prototype implementation over the public Internet are presented as a proof-of-concept.

The remainder of this paper is organized as follows. Section II describes the ASN-FWD mechanism design. Section III analyses ASN-FWD through experimental results. Section IV discusses the adoption of ASN-FWD in the Internet, detailing the possible time window for its deployment. Finally, Section V offers an overview of the related work and Section VI concludes this paper outlining some of our future work.

II. ASN-FWD Design

A. Preliminaries

Current IPv4 traffic forwarding is based on the destination IP address present in the packet, usually included in the header by the source node originating it. Once the packet arrives at a router for processing, the router admits the packet and performs a Longest Prefix Match (LPM) operation based on the destination IP address. Basically, the LPM operation is performed with support of the information stored at the FIB of routers, populated with IP Prefixes reachable all over the Internet. The IP Prefixes are disseminated by ASs mainly using BGP, and each IP Prefix is complemented with the traversed AS-Path information. Firstly, the received BGP information is stored on the RIB (Routing Information Base) of routers, and further processed for FIB fulfilling with the best information according to the routing policies adopted on each AS.

B. Design Goals

The ASN-FWD proposal aims at significantly reducing the number of IPv4 entries on FIB while providing backwards compatibility for existing IPv4 applications. The key idea behind ASN-FWD consists in introducing a single new step (i.e. AS-based forwarding) in the current IPv4 traffic forwarding process briefly described before. The following design goals have guided the ASN-FWD proposal:

1) No changes in the software of routers;
2) No changes in routing protocols currently used in Internet;
3) No changes in the protocol stack of end hosts;
4) No need for stateful middle-boxes for implementing ASN-FWD;
5) No need for additional translation or naming mechanisms;
6) No need for a centralized solution;
7) No dependence on DNS structure, since it has its own problems and challenges;
8) Compatible with current CDNs (Content Delivery Networks);
9) Seamless communication among ASs with and without support to ASN-FWD;
10) Backward compatibility with all IPv4-based applications.

All the discussion about the ASN-FWD design is presented in the following subsections, and is based on the network scenario depicted in Figure 1.

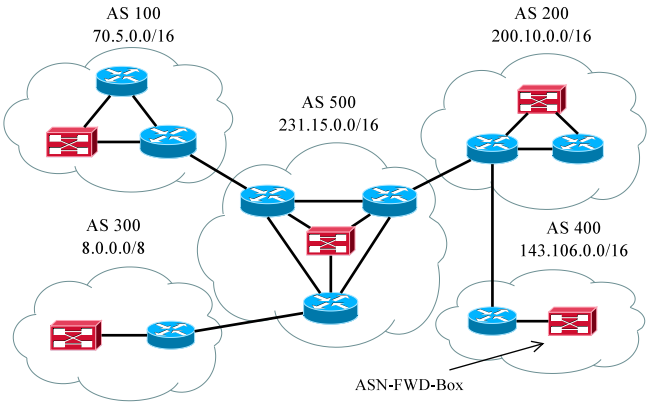

Fig. 1. ASN-FWD exemplification network scenario.

C. Filling the FIB

The main principle of ASN-FWD is allowing IPv4 traffic forwarding based on the 32-bit ASN. To this aim, we now discuss how to insert ASN information in the FIB of routers in a transparent and fully BGP compatible way, requiring no modification on the current software of routers.

Figure 2 brings an example of a simplified BGP message disseminated by AS 500 in the scenario depicted in Figure 1. For simplicity, this message focuses on the Network and AS-Path information, hiding additional information like next hop and weight, also disseminated by BGP. The message shown in Figure 2 represents the current IPv4 inter-domain routing practices in the Internet. Each line contains one IP Prefix available in the network scenario, followed by the correspondent AS-Path. For example, AS 400 disseminates its IP Prefix (143.106.0.0/16) to AS 200. In turn, based on its routing policies, AS 200 disseminates such information to AS 500, informing the AS-Path $200-400$ towards 143.106.0.0/16. Finally, AS 500 disseminates the AS-Path $500-200-400$ towards the IP Prefix available within AS 400. As seen in the highlighted ASNs of Figure 2, the right-most ASN represents the AS responsible for the disseminated IP Prefix.

The basis for the ASN-FWD deployment is utilizing the current 32-bit-long ASN identity space, extended over five years ago from the previous 16-bit-long identity space [10]. Verifying all the current ASNs assigned to IPv4 networks, we find out that none of them presents a value higher than 2^{24}. Leveraging this observation, ASN-FWD proposes to adopt an unused /8 IPv4 Prefix and concatenate such prefix to the 24-less-significant-bits of the assigned ASN of IPv4 networks.

Network	...	AS-Path
143.106.0.0/16		500 200 **400**
70.5.0.0/16		500 **100**
8.0.0.0/8		500 **300**
200.10.0.0/16		500 **200**
231.15.0.0/16		**500**

Fig. 2. Standard BGP information disseminated by AS 500.

For the sake of limiting the scope of this paper, we assume that IANA decides to assign the prefix $11/8$ to be used with ASN-FWD. Consequently, each ASN-FWD-compliant-AS disseminates through BGP its own ASN in the format $11.x.y.z/32$, where $x.y.z$ correspond to the 24 right-most-bits of its own ASN.

Essentially, the main objective behind such concatenation process is to maintain total compatibility with BGP, creating a transparent scenario for including the ASN information in the FIB of routers, requiring no change on the software of routers. Figure 3 presents the simplified BGP message disseminated by AS 500 according to the ASN-FWD proposal. The first five lines represent the ASN dissemination process of ASN-FWD. For example, the ASN 100 is represented by IP Prefix 11.0.0.100/32, reachable through the AS-Path $500 - 100$. The AS 500 is represented by IP Prefix 11.0.1.244/32, where $0.1.244_2$ corresponds to 500_{10}.

Network	...	AS-Path
11.0.0.100/32		500 **100**
11.0.0.200/32		500 **200**
11.0.1.44/32		500 **300**
11.0.1.144/32		500 200 **400**
11.0.1.244/32		**500**
143.106.0.0/16		500 200 400
70.5.0.0/16		500 100
8.0.0.0/8		500 300
200.10.0.0/16		500 200
231.15.0.0/16		500

Fig. 3. ASN-FWD-extended BGP information disseminated by AS 500.

Despite the ASN information being transparently disseminated via BGP, in order to fully support ASN-FWD, there is a requirement for policy changes due to the fact that, currently, most of the carriers are configured to filter $/32$ prefixes. In this way, carriers need to create a routing policy allowing $/32$ advertisements via E-BGP for prefixes beginning with $11/8$.

D. Preparing IPv4 Packets for ASN-FWD Operation

In essence, the IPv4 operation of ASN-FWD relies on adding a 32-bits-long-ASN information per packet, representing the ASN from which the IP Prefix (the LPM match), correspondent to the destination IP marked in the packet, was disseminated with BGP. One way to introduce the ASN information on the packet headers is to use the optional headers of IPv4, thus adding extra 8-bytes per packet. Figure 4 details the ASN-FWD header format.

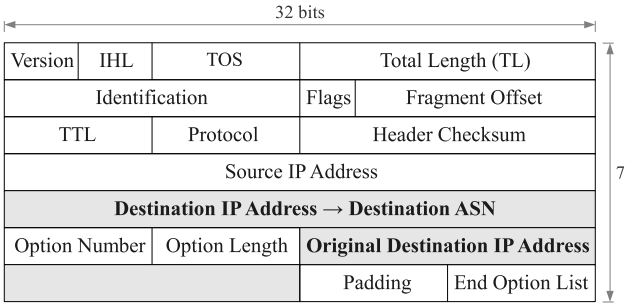

Fig. 4. ASN-FWD IPv4 packet header format.

When compared to the legacy IPv4 forwarding process, the ASN-FWD mechanism introduces a new step in the overall IPv4 traffic handling process. The resulting process modifies the packet (*adaptation*) header at the network layer based on the mapping from the original Destination IP Address to the correspondent Destination ASN information. Such an IP-to-ASN query is already available in [11],[1] and can be easily implemented by available boxes that perform, for example, CGN (Carrier Grade NAT) [12] modification on packets. Alternatively, this programmable forwarding pipeline can be also implemented following a Software Defined Networking (SDN) [13] approach, for example, implemented using OpenFlow [14] match and action instructions. The details for SDN implementation options fall beyond the scope of this paper. [2]

In order to be transparent to IPv4 legacy nodes and applications, ASN-FWD adopts a network-centric proposal, including ASN-FWD-Boxes on a per AS basis for *packets adaptation*. As an alternative, future mechanisms can be moved to a host-centric solution, delegating to IPv4 nodes the responsibility of generating packets according to the proposed format as shown in Figure 4.

Besides the packets adaptation, the IPv4 forwarding mechanism in routers remains unchanged, i.e., it is still based on the destination IP address present in the header, except to the fact that the forwarding actions are actually taken based on the ASN information included by the ASN-FWD-Boxes. The overall ASN-FWD packet adaptation process such as inserting and removing the optional header is detailed in Figure 5.

To exemplify a packet transmission in the ASN-FWD mechanism, consider a source host with IP 70.5.0.1 located at AS 100, which starts a communication with a destination host located at AS 400, whose IP is 143.106.0.1. The source node generates a packet with source IP address 70.5.0.1 and destination IP address 143.106.0.1 and injects it on the Internet. Its AS 100 directs the packet to one of the available ASN-FWD-Boxes, that rewrites the packet according to the process described in Figure 5. As a result, the packet is

[1]The service available in [11] needs to be adapted to return the ASN information in the proposed format $11.x.y.z/32$.

[2]Related efforts on SDN Mapping Service using the OpenDaylight controller can be found here: http://wiki.opendaylight.org/view/OpenDaylight_Lisp_Flow_Mapping

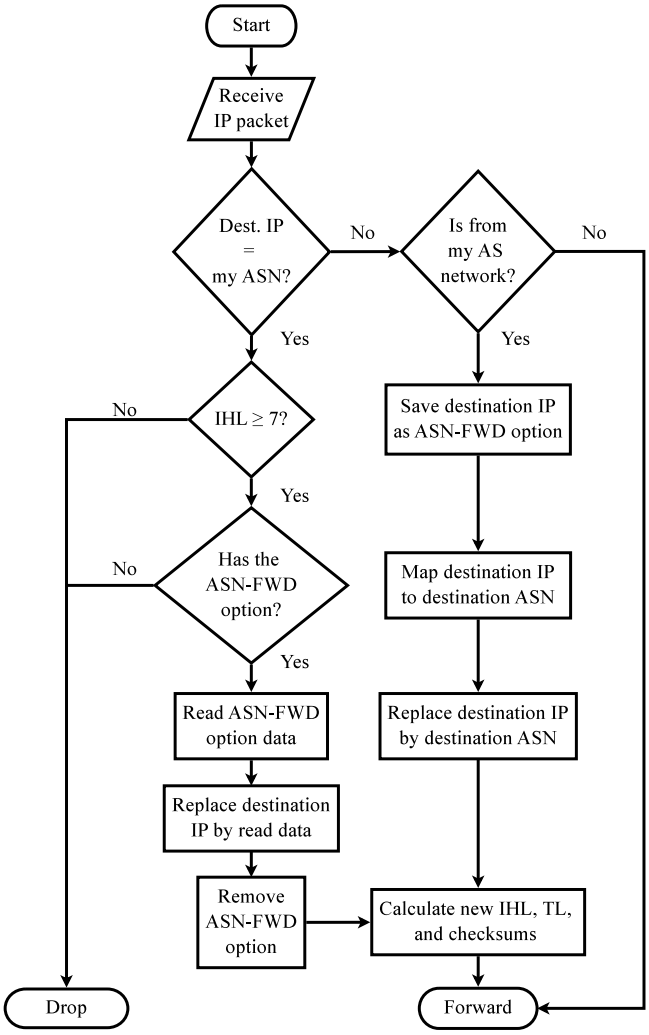

Fig. 5. Flowchart of ASN-FWD adaptation of packets.

forwarded throughout the Internet with source IP address 70.5.0.1, destination IP address 11.0.1.144 (400_{10}), and the IP address 143.106.0.1 written in the optional header field.

In order to reach the destination AS 400, the proposed mechanism relies in the presence of information regarding the ASNs in the FIBs, generated according to the process detailed in Section II-C. Note that once the packet is modified to the ASN-FWD format, the forwarding process by all routers in the end-to-end path is exactly the same as currently performed, i.e., no change are required in the routers in order to support ASN-FWD.

Once the packet is received by the destination AS 400, it is pushed towards one of its ASN-FWD-Box, and is rewritten according to the process of Figure 5, recovering the original IPv4 format, i.e., the destination IP address is rewritten to 143.106.0.1. After the rewriting process, the packet is locally delivered to node 143.106.0.1 and traverses its legacy TCP/IP stack up to the application, receives the required treatment, and the answer is sent back to the host 70.5.0.1.

The legacy answer packet from 143.106.0.1 to 70.5.0.1 is pushed across one of the ASN-FWD-Boxes located at AS 400, which based on the destination IP 70.5.0.1 obtains the corresponding AS 100 and rewrites the packet according to the process described in Figure 5 - Source IP address 143.106.0.1, destination IP address 11.0.0.100, and the IP address 70.5.0.1 written in the optional header field. Next we will present TCP and UDP experiments performed with legacy applications that demonstrate the ASN-FWD operation, following the process described in this example.

III. ASN-FWD Experimental Implementation

We turn our attention now to an experimental evaluation of ASN-FWD, performed with a prototype developed using *libipq* [15]. The objective is to validate the ASN-FWD-Box, transmitting IPv4 packets through the public Internet according to the ASN-FWD header format. The experiments considered TCP and UDP transmissions, using the legacy *wget* and *netcat* applications, respectively, and were performed connecting two networks located in Brazil. The first network was set in the Federal University of Uberlândia, connected to the public Internet through the RNP backbone [16]. In this network, it was used the IP 200.19.151.21 as the ASN, and the pool 200.19.151.32/30 as an IP Prefix reachable from such ASN. The second network was set in a Brazilian's operator, in which a single host was set at the IP 189.15.69.57, i.e., this IP corresponds to the ASN and also as an IP Prefix reachable from there.

At the university side, the prototype runs in an Intel Core2 Quad CPU Q9550 2.83GHz with 4GB of RAM. In the operator side, the prototype runs in an Intel Core i7-2640M 2.8GHz with 6GB of RAM. Both machines use OpenSuSE 12.2 with Linux (kernel 3.4.47), in which VirtualBox 4.2.12 provides the virtual machines running the end hosts and the ASN-FWD-Boxes. Figure 6 presents the *traceroute* of all experiments, detailing the routers connecting the source to destination nodes, all of them transparently forwarding the ASN-FWD packets.

```
traceroute to 189.15.69.171 (189.15.69.171), 30 hops max, 40 byte packets
using UDP
 1  200.19.151.254 (200.19.151.254)  0.000 ms   0.000 ms   0.000 ms
 2  * * *
 3  200.131.199.121 (200.131.199.121)  8.866 ms   7.459 ms   15.259 ms
 4  200.131.199.17 (200.131.199.17)  2.166 ms   1.991 ms   2.601 ms
 5  200.19.158.245 (200.19.158.245)  23.721 ms   22.730 ms   21.959 ms
 6  200.143.255.173 (200.143.255.173)  22.480 ms   21.363 ms   20.187 ms
 7  200.219.139.109 (200.219.139.109)  33.140 ms   31.978 ms   39.497 ms
 8  201.48.45.165 (201.48.45.165)  39.677 ms   37.577 ms   37.441 ms
 9  * * *
10  189.15.69.171 (189.15.69.171)  51.785 ms   55.920 ms   67.576 ms
```

Fig. 6. End-to-end path of the experiments collected with *traceroute*.

The first experiment uses the legacy TCP application *wget*. A web server is set in a virtual machine using the IP 200.19.151.34 inside the 200.19.151.32/30 network at the university, represented by the ASN 200.19.151.21. The *wget* is issued at the operator's network from the IP 189.15.69.171. Figure 7 shows the communication log for this experiment. As can be seen in the top of Figure 7(a), the source/destination of packets does not match. In the operator to the university

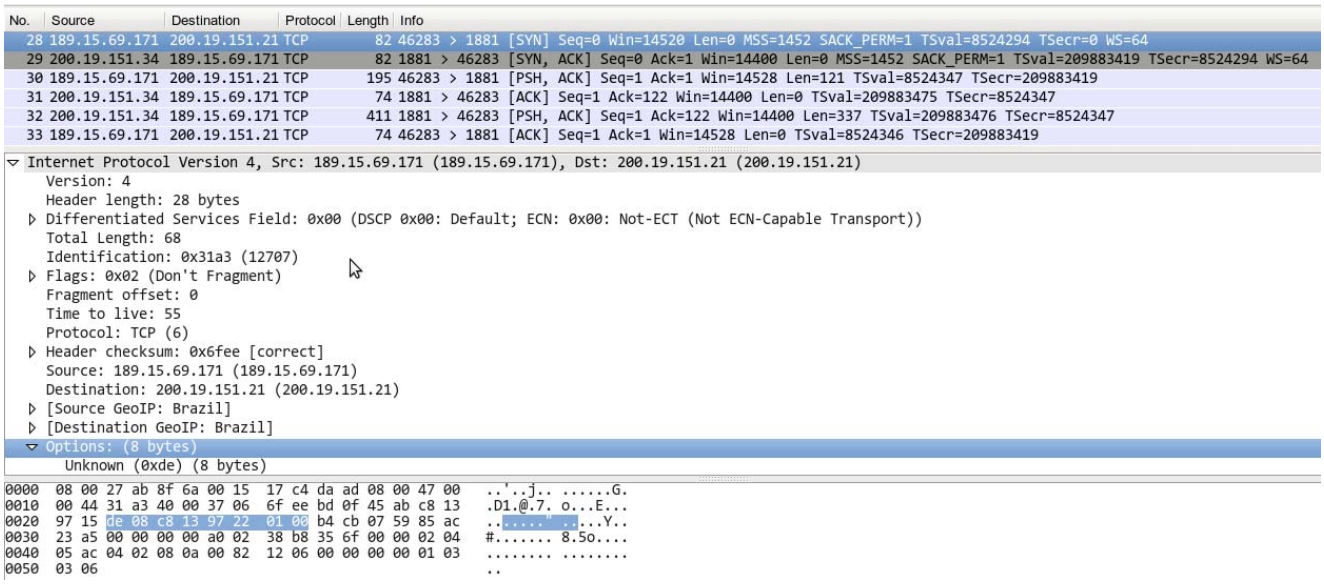

(a) From the operator to the university.

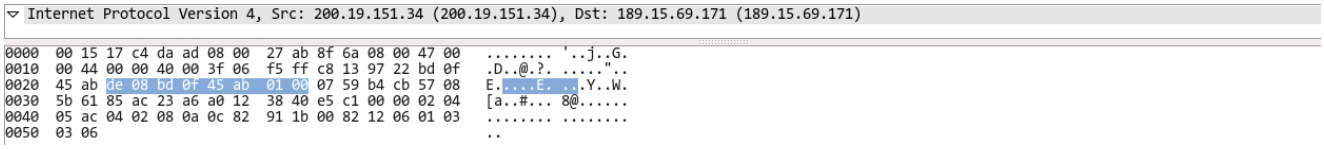

(b) From the university to the operator.

Fig. 7. ASN-FWD TCP log collected using the legacy $wget$ application.

direction, the source IP is 189.15.69.171 (end-host IP) and the destination IP is 200.19.151.21 (ASN). On the other hand, the source IP is 200.19.151.34 (end-host IP) and the destination IP is 189.15.69.171 (ASN).

Figure 7(a) details the IP layer information of a packet, from which it is possible to verify the existence of an Optional Header of 8-bytes length using the experimental type $0xDE$ (222_{10}). In the bottom part of Figure 7(a), it is possible to find the hexadecimal representation of the optional header highlighted in blue. Note the presence of the information $c8$ 13 97 22, which corresponds to 200.19.151.34, the original destination IP. In essence, this packet was firstly generated by the legacy node running the $wget$ application, using source IP 189.15.69.171 and destination IP 200.19.151.34. In the sequence, this packet was pushed towards the ASN-FWD-Box, which moved the original destination IP 200.19.151.34 to the optional header, and based on the IP-to-ASN mapping, wrote the destination IP as 200.19.151.21.

When the packet arrives at the ASN-FWD-Box running in 200.19.151.21, it is restored to the legacy format and delivered to the web server running in the IP 200.19.151.34. The request is transparently handled and the answer packet is sent back to the requesting node. The source IP of the answer is 200.19.151.34 and the destination IP is 189.15.69.171. This legacy packet is pushed to the ASN-FWD-Box 200.19.151.21, which adapts the packet to source IP 200.19.151.34, destination IP 189.15.69.171 (ASN) and the optional header carries the original destination IP 189.15.69.171, depicted in Figure 7(b) by the bd $0f$ 45 ab highlighted hexadecimal information.

The second experiment uses the legacy application $netcat$ with UDP, allowing the user to exchange messages through a console. As in the previous experiments, at the university side the $netcat$ is executed in the virtual machine with IP 200.19.151.34, represented by the ASN 200.19.151.21, and on the operator side the virtual machine has the IP 189.15.69.171 (end-host IP and ASN). Figure 8 shows the communication log for this experiment. As can be seen in the top of Figure 8(a), once again the source/destination of exchanged packets are different.

At the bottom part of Figure 8(a), it is possible to find the hexadecimal representation of the optional header, highlighted in blue. Note the presence of the information $c8$ 13 97 22, which corresponds to 200.19.151.34, the original destination IP. On the opposite direction, Figure 8(b) shows in the optional packet header the information bd $0f$ 45 ab, correspondent to 189.15.69.171. As a future work, international communications through the public Internet will be used to analyze the ASN-FWD transmissions.

As shown in Figures 7 and 8, the changes required by ASN-FWD at the IP header incur extra 8-bytes per packet. To conclude the ASN-FWD analysis, Figure 9 plots the per

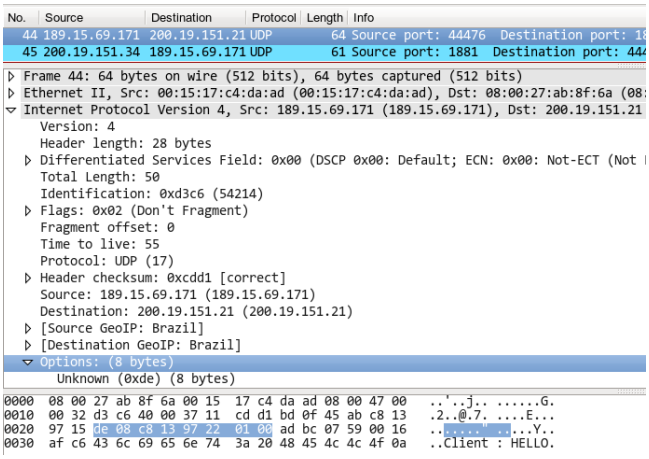

(a) From the operator to the university.

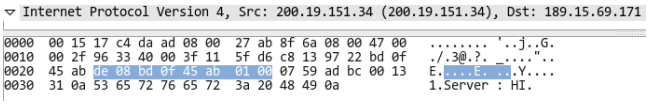

(b) From the university to the operator.

Fig. 8. ASN-FWD UDP log collected using the legacy *netcat* application.

packet overhead incurred by the ASN-FWD proposal as a function of the amount of bytes to be transmitted. This figure considers two common MTU (Maximum Transfer Units) sizes, 576 and 1500 bytes. As observed in this figure the overhead is negligible, specially for the most common MTU of 1500 bytes.

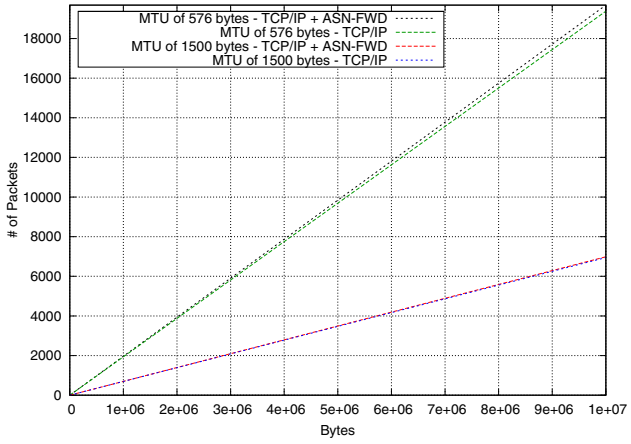

Fig. 9. Packet overhead of ASN-FWD in TCP/IP transmissions.

One important aspect related to the MTU size is to try to avoid IPv4 packets segmentation. In this way, once the ASN-FWD-Box receives the first communication packets from the legacy nodes, the ASN-FWD-Box sends an ICMP Too Big message back to the source node, informing the best MTU size, defined by the effective MTU size minus the 8-Bytes of the ASN-FWD proposal.

IV. On the ASN-FWD Deployment Time Window

The analysis of ASN-FWD deployment time window is two fold: (*i*) the current capacity of high capacity routers for storing entries on FIB and (*ii*) the FIB evolution during the migration from IPv4 to IPv6 being faced by the Internet.

In respect to the first aspect, most of the high capacity routers being used nowadays support one million IPv4 entries on FIB, sharing such space between IPv4 and IPv6 entries. By default, most of the vendors sell routers on a fifty/fifty configuration and carriers usually maintain such configuration, allowing the storage of 512K IPv4 prefixes and 256K IPv6 prefixes [17]. An IPv6 FIB entry presents the double of the size of an IPv4 FIB entry.

From the technical perspective, evolving FIB capacity can lead vendors to important challenges, including not only memory and processing speeds, but also power and heat dissipation. From the business perspective, evolving FIB capacity can push carriers to frequent boxes upgrades, increasing CAPEX and OPEX, what might lead carriers to take some decisions for cutting their costs. For example, there are smaller ISPs nowadays that, instead of upgrading their boxes, implement routing policies for filtering $/24$ prefixes, which implies that parts of the Internet are not reachable [6] from such ISPs.

In respect to the FIB evolution during the transition phase, Figure 10 starts by presenting the historical growth of the number of IPv4 and IPv6 entries on FIB, extracted from the BGP Reports [9]. As can be seen in this figure, the current number of IPv4 entries on FIB is around 480K and the current number of IPv6 entries is around 15K. From this information, there is a first impact on the current routers configuration, once the factory default 512K IPv4 number of entries are practically reached, forcing ISPs to change it, for example, to 768K IPv4 entries and 128K IPv6 entries, in order to continue their normal operation on the upcoming years. Based on such values, the upper bound of the y axis of Figure 10 is 768K entries, and there is a reference line plotted at the 128K entries value to guide the deployment time window discussion.

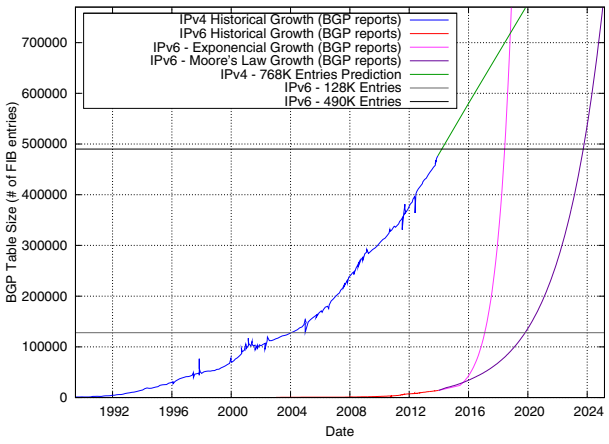

Fig. 10. ASN-FWD Deployment Time Window Analyses.

The next aspect is to predict the size of IPv6 in the near future, which is a very difficult task. Analysing facts available in the literature, there are some misbehaves already identified by the IETF Softwire group [18], conducting the IPv4 to IPv6 transition phase. For example, the ISPs are performing an one-to-one IPv4 to IPv6 migration, i.e., for each existent IPv4 prefix, ISPs are assign an IPv6 prefix. The ideal methodology, according to the IETF Softwire group, is to adopt IPv6 as a mechanism for recovering the strong prefix aggregation nature of Internet routing, avoiding an explosive growth in the number of IPv6 entries on FIB. Another fact available in the literature, as seen in the historical IPv6 growth depicted in Figure 10, is the slowness in the IPv6 adoption, not even boosted with the worldwide official IPv6 launch day on 2011 [19]. The reasons for such slowness include, for example, lack of financial investment and lack of IPv6 culture.

In this challenging context, BGP-Reports present two IPv6 growing analyses, also included in Figure 10, the first one exhibiting an exponential growth and the second one exhibiting a growing rate according to the Moore's Law [9]. Based on the observed slowness of IPv6 adoption and in the careful intervention of IETF Softwire group, we consider in this paper the Moore's Law prediction for building the ASN-FWD deployment time window discussion presented in this section.

Observing the Moore's Law plot in Figure 10, the upper bound of 128K IPv6 entries might be achieved in 2020, providing about to 6 years for the ASN-FWD deployment. To exemplify how the number of IPv4 entries should evolve during this time window, in order to achieve the 768K upper bound, Figure 10 plots a prediction line. Such prediction presents a strong growing rate, opposite to the growing trend pointed by BGP-Reports [9], which indicates a reduction in the pace since the depletion of the IPv4 public pool of addresses. Consequently, we consider that the deployment time window can be bigger than 6 years.

After the deployment of ASN-FWD, the IPv4 share on FIB is reduced to the number of IPv4 ASNs, currently around 45K ASs (approximately 10% of its current share), paving the way for the IPv6 growth up to the limit of roughly 490K entries. As seen in the 490K reference line plotted in Figure 10, it might occur in 2024 according to the Moore Law's trend, providing at least a decade of pressure reduction over the technological and business pressures on the Internet.

V. Related Work

A myriad of efforts have been devoted to addressing the scalability issues of the global Internet, the pressure on IPv4 exhaustion, the migration to IPv6, and additional research topics on related Internet routing and addressing. In the following, we discuss an incomplete set of related work with similarities and intellectual contributions to ASN-FWD.

Locator/Identifier Separation Protocol (LISP): The origins of LISP can be rooted back to discussions during the IAB-sponsored Routing and Addressing Workshop held in late 2006 [20]. To address the identified scaling issues, the approach proposed by LISP is to replace IP addresses with two new types of numbers: (*i*) topologically assigned Routing Locators (RLOCs) used for routing and forwarding of packets through the network; and (*ii*) not globally routable Endpoint Identifiers (EIDs), which are topology-independent device identifiers aggregated along administrative boundaries. The LISP specification [5] defines functions for mapping between the two numbering spaces and for encapsulating traffic originated by devices using non-routable EIDs for transport across a network infrastructure that routes and forwards using RLOCs. Both RLOCs and EIDs are syntactically identical to IP addresses but fundamentally different in their usage semantics.

Host Identity Protocol (HIP): The HIP architecture [21] proposes an alternative to the dual use of IP addresses as "locators" (routing labels) and "identifiers" (endpoint/host identifiers). In HIP, public/private key pairs are used as Host Identifiers, to which higher layer protocols are bound instead of an IP address. By using public keys (and their representations) as host identifiers, dynamic changes to IP address sets can be directly authenticated between hosts. Thereby, HIP introduces a new namespace (Host Identity namespace) and only requires changes at endpoints (in addition to infrastructure services such as Rendezvous Server). However, this id-loc split does not help in reducing the pressure on the global routing infrastructure.

A Better Internet Without IP Addresses: Routing based in ASN has been proposed before in [8], in which, the IP stack is modified to eliminate IP addresses, where hosts would be located by their names. Therefore, host names would be used in the packet header for forwarding inside an AS, and, for allowing scalability, inter-AS forwarding would be performed based on the AS number, also carried in the packet headers. One caveat of this proposal is requiring many changes to the deployed Internet infrastructure, including end hosts and routers.

ViAggre: Virtual Aggregation [6] is a configuration-only approach to shrink the routing table on routers. It can be independently and autonomously adopted by ISPs, and is presented as a short-term solution to slow down the increase of DFZ routing tables, while any long-term architectural change hasn't been agreed. ViAggre divides the global addressing space into a set of virtual prefixes, larger than any aggregated prefix in use, and assigns each virtual prefix to a certain router inside a given ISP. In this way, routers do not need to maintain information regarding all IP Prefixes reachable in the Internet. Once a packet is received in such ViAggre ISP, it is tunneled from the router that received the packet to the router responsible for its correspondent IP Prefix, and in the sequence from this router towards the router responsible for the external connection to the next hop.

SMALTA: Saving Memory And Lookup Time via Aggregation [7] is a FIB aggregation scheme that reduces both FIB size and lookup time, at the cost of increasing delay for inserting updates into the FIB. In a standard scenario, the route resolution function receives information from multiple sources,

such as BGP, and generates a table of best-match prefixes and equivalent next hops. SMALTA takes this standard FIB table and generates a new one with fewer entries. This can be achieved by finding pairs of entries that have same next hops and can be reduced to a single line. For instance, if the entries 2.0.0.0/8 and 3.0.0.0/8 are present in the table, they can be aggregated into 2.0.0.0/7, as long as they have the same next hop. The new aggregated tree is obtained by an incremental FIB aggregation scheme based on Optimal Routing Table Constructor (ORTC) [22]. The authors assure that SMALTA can reduce FIB storage size by at least 50%, but besides the delay for generating the FIB, mechanisms like SMALTA make difficult online FIB updates, since information must be pre-aggregated in order to avoid inconsistencies at FIB.

VI. Conclusions and Future Work

ASN-FWD shrinks the IPv4 share on FIB and offers backwards compatibility to IPv4 legacy applications. It is a minimally invasive proposal, developed transparently on top of standardized mechanisms, such as optional IP headers, BGP messages, FIB generation functions and packets' forwarding mechanisms. The only new requirement is the IP-to-ASN mapping, also developed on top of information already available through BGP. The main contribution of ASN-FWD is the ASN-based forwarding of information in the Internet, that can also be adopted for IPv6. For the IPv4 enthusiasts, ASN-FWD breaks the IPv4 exhaustion, since nodes are uniquely identified by the tuple $ASN + IP$, i.e. for each ASN in the Internet, a complete 32-bits IPv4 address space can be allocated.

Regarding the future work, at the current stage ASN-FWD opens up some possibilities. Firstly, the ASN-FWD prototype needs to receive some extensions in order to be packed and made available for download. Such distribution would allow a deeper investigation on ASN-FWD operation, counting on the involvement of volunteers from the worldwide research community. Another important aspect to be investigated is the ASN-FWD-Box deployment using SDN mechanisms (e.g. OpenFlow match + encap/re-write actions), an approach that can ease the adoption of ASN-FWD. In addition, host-based approaches in data center scenarios will be investigated.

Acknowledgment

The authors would like to thank CAPES for supporting this work.

References

[1] T. Bates, P. Smith, and G. Huston, "The CIDR Report," 2013. [Online]. Available: http://www.cidr-report.org

[2] N. Stringfield, R. White, and S. McKee, *Cisco Express Forwarding (Networking Technology)*. Cisco Press, 2007.

[3] T. Li, "Design Goals for Scalable Internet Routing," *IETF RFC-6227*, May 2011.

[4] R. Moskowitz and P. Nikander, "Host Identity Protocol (HIP) Architecture," *IETF RFC-4423*, May 2006.

[5] D. Farinacci, V. Fuller, D. Meyer, and D. Lewis, "Locator/ID Separation Protocol (LISP)," *IETF RFC-6830*, January 2013.

[6] H. Ballani, P. Francis, T. Cao, and J. Wang, "Making routers last longer with viaggre," in *Proceedings of the 6th USENIX symposium on Networked systems design and implementation*, ser. NSDI'09. Berkeley, CA, USA: USENIX Association, 2009, pp. 453–466. [Online]. Available: http://dl.acm.org/citation.cfm?id=1558977.1559008

[7] Z. A. Uzmi, M. Nebel, A. Tariq, S. Jawad, R. Chen, A. Shaikh, J. Wang, and P. Francis, "Smalta: practical and near-optimal fib aggregation," in *Proceedings of the Seventh COnference on emerging Networking EXperiments and Technologies*, ser. CoNEXT '11. New York, NY, USA: ACM, 2011, pp. 29:1–29:12. [Online]. Available: http://doi.acm.org/10.1145/2079296.2079325

[8] C. Shue, "A better internet without ip addresses," Ph.D. dissertation, Indianapolis, IN, USA, 2009, aAI3358945.

[9] G. Huston, "The BGP Report," 2012. [Online]. Available: http://conference.apnic.net/34/home

[10] Q. Vohra and E. Chen, "BGP Support for Four-octet AS Number Space," *IETF RFC-4893*, May 2007.

[11] T. Cymru, "IP to ASN Mapping," 2013. [Online]. Available: http://www.team-cymru.org/Services/ip-to-asn.html

[12] S. Perreault, I. Yamagata, S. Miyakawa, A. Nakagawa, and H. Ashida, "Common Requirements for Carrier-Grade NATs (CGNs)," *IETF RFC-6888*, April 2013.

[13] M. Mendonça, B. N. Astuto, X. N. Nguyen, K. Obraczka, and T. Turletti, "A Survey of Software-Defined Networking: Past, Present, and Future of Programmable Networks," Jun. 2013, in Submission In Submission. [Online]. Available: http://hal.inria.fr/hal-00825087

[14] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "Openflow: enabling innovation in campus networks," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 69–74, Mar. 2008. [Online]. Available: http://doi.acm.org/10.1145/1355734.1355746

[15] "The netfilter.org 'iptables' project," 2013. [Online]. Available: http://www.netfilter.org/

[16] "Rede Nacional de Ensino e Pesquisa," 2013. [Online]. Available: http://www.rnp.br

[17] "IPv4/IPv6 and TCAM memory," 2012. [Online]. Available: http://www.ipv4depletion.com/?p=672

[18] "The IETF Softwire Working Group," 2013. [Online]. Available: http://datatracker.ietf.org/wg/softwire/charter/

[19] "World IPv6 Launch," 2011. [Online]. Available: http://www.worldipv6launch.org/

[20] D. Meyer, L. Zhang, and K. Fall, "Report from the IAB Workshop on Routing and Addressing," RFC 4984 (Informational), Internet Engineering Task Force, Sep. 2007. [Online]. Available: http://www.ietf.org/rfc/rfc4984.txt

[21] R. Moskowitz and P. Nikander, "Host Identity Protocol (HIP) Architecture," RFC 4423 (Informational), Internet Engineering Task Force, May 2006. [Online]. Available: http://www.ietf.org/rfc/rfc4423.txt

[22] R. P. Draves, C. King, S. Venkatachary, and B. D. Zill, "Constructing optimal ip routing tables," in *INFOCOM'99. Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 1. IEEE, 1999, pp. 88–97.