A Brief History of MPLS Usage in IPv6

Yves Vanaubel[1], Pascal Mérindol[2], Jean-Jacques Pansiot[2], Benoit Donnet[1] ⋆

[1] Université de Liège – Belgium,
[2] Université de Strasbourg – France

Abstract. Recent researches have stated the fast deployment of IPv6. It has been demonstrated that IPv6 grows much faster, being so more and more adopted by both Internet service providers but also by servers and end-hosts. In parallel, researches have been conducted to discover and assess the usage of MPLS tunnels. Indeed, recent developments in the ICMP protocol make certain categories of MPLS tunnels transparent to `traceroute` probing. However, these studies focus only on IPv4, where MPLS is strongly deployed.
In this paper, we provide a first look at how MPLS is used under IPv6 networks using `traceroute` data collected by CAIDA. At first glance, we observe that the MPLS deployment and usage seem to greatly differ between IPv4 and IPv6, in particular in the way MPLS label stacks are used. While label stacks with at least two labels are marginal in IPv4 (and mostly correspond to a VPN usage), they are prevalent in IPv6. After a deeper analysis of the label stack typical content in IPv6, we show that such tunnels result from the use of 6PE. This is not really surprising since this mechanism was specifically designed to forward IPv6 traffic using MPLS tunnels through networks that are not fully IPv6 compliant. However, we show that it does not result from non dual-stack routers but rather from the absence of native IPv6 MPLS signaling protocols. Finally, we investigate a large Tier-1 network, Cogent, that stands out with an original set-up.

Keywords: IPv6 · 6PE · network discovery · MPLS · LDP · RSVP-TE · traceroute

1 Introduction

During the last years, IPv6 has drawn the attention of the research community. For instance, Dhamdere et al. [1] showed that IPv6 is differently deployed over the world (IPv6 is more present in Europe than in the USA), while the routing dynamics and the path performance are largely identical between IPv4 and IPv6. More recently, Czyz et al. [2] showed that IPv6 networks are becoming mature and entering now a production mode. Further, on September, 24$^{\text{th}}$, 2015, the ARIN IPv4 free pool reached zero, effectively triggering full IPv4 depletion. The ARIN is now unable to provide any IPv4 block except for those requiring

⋆ This work is partially funded by the European Commission funded mPlane ICT-318627 project.

a small block in order to ease the IPv6 transition [3]. We believe this should accelerate the global IPv6 adoption.

In parallel to this IPv6 interest, MPLS has been more and more investigated by the research community. For instance, Sommers et al. [4] examined the characteristics of MPLS deployments that are explicitly identified using RFC4950 extensions. Donnet et al. [5] provided algorithms for detecting MPLS tunnels depending on the way MPLS routers react to the `ttl-propagate` and RFC4950 options. Others looked at the MPLS usage. Pathak et al. [6] quantified the additional delay caused by MPLS when used for traffic engineering (TE) reasons. More recently, Vanaubel et al. [7] evaluated the MPLS usage in the light of transit path diversity, showing that the basic usage for scalability purpose (e.g., with LDP) seems predominant, with or without path diversity and that TE is well represented in a subset of specific ASes. None of those works investigated MPLS under IPv6.

As the deployment of IPv6 is growing and the interest in MPLS is stronger, we aim, in this paper, to investigate the state of MPLS deployment under IPv6. In particular, we are interested in knowing how operators are using MPLS in IPv6 and whether this usage differs from the one in IPv4. To achieve this goal, we rely on an IPv6 `traceroute` dataset collected by CAIDA between 2009 and 2015. From this dataset, we extract tunnels [5] and show that, in parallel to an increase in the IPv6 deployment, there is, along the time, an increase in the MPLS usage in IPv6. This usage, as we show it latter in the paper, is essentially oriented for 6PE purpose (i.e., either for connecting IPv6 islands together or using LDP for IPv4 to build tunnels carrying both IPv6 and IPv4 traffic on dual stack MPLS routers). We also investigate the particular case of Cogent, a large Tier-1 ISP having both a very prominent position in the dataset and a very particular behavior in regards to 6PE.

The remainder of this paper is organized as follows: Sec. 2 provides the required background for this paper. Sec. 3 presents our findings. Finally, Sec. 4 concludes this paper by summarizing its main achievements.

2 Background

2.1 MPLS Overview

The *Multiprotocol Label Switching* (MPLS) [8] was originally designed to speed up the forwarding process. In practice, this was done with one or more 32 bits *label stack entries* (LSE) inserted between the frame header (Data-link layer) and the IP packet (Network layer). A given packet may carry out several LSEs at the same time. In this case, the packet is said having a *stack of labels*. Each LSE is made of four fields: a 20-bit label value used for forwarding the packet to the next router, a 3-bit Traffic Class field for quality of service (QoS), priority, and Explicit Congestion Notification (ECN) [9], a 1-bit bottom of stack flag (when set the current label is the last in the stack [10]), and an 8-bit time-to-live (LSE-TTL) field having the same purpose as the IP-TTL field [11].

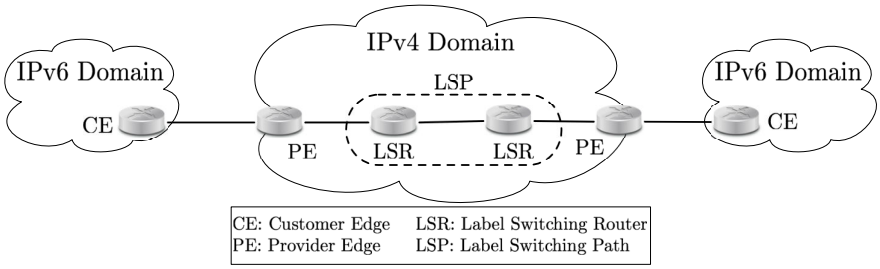

Fig. 1. 6PE usage of MPLS. PE routers are dual-stack, while LSRs are IPv4 only routers.

MPLS routers, called *Label Switching Routers* (LSRs), exchange labelled packets over *Label Switched Paths* (LSPs). The first MPLS router (*Ingress Label Edge Router*, or Ingress LER, i.e., the tunnel entry point) adds the label stack, while the last MPLS router (*Egress Label Edge Router*, or Egress LER, i.e., the tunnel exit point) removes the label stack. In some cases, and in particular with Cisco routers, the LSE stack may be removed by the penultimate MPLS router (*penultimate hop popping*, PHP) to reduce the MPLS overhead. The Egress LER then performs a classic IP lookup and forwards the traffic, reducing so the load on the Egress LER (especially if the Egress LER is shared among several LSPs). This means that, when using PHP, technically speaking, the MPLS tunnel exit is one hop before the Egress LER. In its most basic operation, LSPs are constructed along best effort routes using the *Label Distribution Protocol* (LDP [12]). More specific LSPs may be constructed for Traffic Engineering purposes, using an extension of the RSVP protocol, *RSVP-TE* [13]. In these two cases, the label stack contains only one LSE. A more complex usage is for Virtual Private Networks (VPN [14]), where LSPs are constructed using either LDP or RSVP-TE, and an additional LSE at the bottom of the label stack is used to specify a Virtual Routing Table at the Egress. In this case, the bottom of the stack is constant along an LSP, while the top of the stack is modified at each hop, as in the previous cases.

2.2 MPLS in IPv6

MPLS can be used in IPv6-only networks in the same way as it is used in IPv4 networks (see George and Pignatoro [15] for a discussion on gaps remaining between IPv4 and IPv6). Indeed most routing protocols and label distribution protocols [16, 17] have now their IPv6 version. However this has not always been the case. Moreover, providers do not activate IPv6 capabilities even when they are available in their hardware and software. Therefore, specific mechanisms have been devised to deliver IPv6 traffic across networks where there is either no IPv6 routing (IPv4 only networks) or where some mechanisms are not IPv6-aware such as LDP [16, 12].

Thus, one of the MPLS usage under IPv6 is to connect IPv6 islands through an IPv4 core network that is unaware of IPv6. This mechanism is called *6PE* [18]

and is illustrated in Fig. 1. This is done through the usage of *Provider Edge* (PE) routers that are dual-stack and that are located at the edge of the IPv4 domain. Each PE router receives IPv6 prefixes from the *Customer Edge* (CE) router in the IPv6 domain. IPv6 reachability is exchanged between 6PEs via multiprotocol-iBGP, MP-BGP.

When 6PE was released, the main objective was to ensure IPv6 connectivity on top of MPLS core routers that are not IPv6-aware. That situation drove the need for two labels in the data plane (due to the potential usage of PHP in particular): (*i*), the top label is the transport label, which is assigned hop-by-hop [12, 13] and, (*ii*), the bottom label is a label assigned by BGP and advertised by iBGP between the PE routers. Quoting RFC4798 [18], "This label advertised by the egress 6PE Router with MP-BGP MAY be an arbitrary label value, which identifies an IPv6 routing context or outgoing interface to send the packet to, or MAY be the IPv6 Explicit Null Label". This last label has a value of 2 [10].

In that context, the PE routers that perform 6PE are the Ingress and Egress LERs. Note that today, now that global IPv6 deployment is more common, 6PE is also interesting for core LSRs with dual-stack routers and IPv6 connectivity. This is useful to build LSP for IPv6 without using an IPv6 label distribution protocol (LDP for IPv6 [16] has been finalized only recently), and/or for sharing the same LSP for IPv4 and IPv6 traffic, reducing so the control plane churn. Our analysis will show that this specific behavior is the most common in practice.

2.3 Revealing MPLS Tunnels

MPLS routers may send ICMP `time-exceeded` messages when the LSE-TTL expires (in both IPv4 and IPv6). In order to debug networks where MPLS is deployed, routers may also implement RFC4950 [19], an extension to ICMP allowing a router to embed an MPLS LSE in an ICMP `time-exceeded` message. In that case, the router simply quotes the MPLS LSE (or the LSE stack) of the received packet in the ICMP `time-exceeded` message. RFC4950 is particularly useful for operators as it allows them to verify the correctness of their MPLS tunnels and TE policy.

If the Ingress LER copies the IP-TTL value to the LSE-TTL field rather than setting the LSE-TTL to an arbitrary value such as 255, LSRs along the LSP will reveal themselves when using traceroute via ICMP messages even if they do not implement RFC4950. Operators can configure this action using the `ttl-propagate` option provided by the router manufacturer [11] (while, to the best of our knowledge, the RFC4950 is just a matter of implementation and cannot be deactivated on recent routers supporting it). These mechanisms are identical for IPv4 and IPv6.

In this paper we focus on *explicit MPLS tunnels*, i.e., tunnels that can be fully revealed via `traceroute` as they implement both TTL propagation (they are seen in traces) and RFC4950 (they are seen as LSRs providing their LSE). Note that in the case of 6PE, if the TTL of a `traceroute` packet expires inside the IPv4 core, the IPv4 router may be unable to send an ICMPv6 error message. In this case, the `traceroute` will be incomplete and the non-responding hop will

Year	Probing				Addresses			Tunnels	
	VPs	Traces	Prefixes	ASes	v6	v4 map'd v6	MPLS	Raw	Complete
2009	5	7,765	2,128	988	4,009	0	14	47	68%
2010	8	17,472	3,550	1,363	6,331	21	48	59	52%
2011	13	51,636	8,347	2,365	12,307	211	199	1,235	22%
2012	21	154,791	18,589	3,918	23,225	704	680	2,783	42%
2013	25	256,725	25,891	4,992	33,239	370	1,468	14,366	45%
2014	29	772,461	32,391	6,224	43,309	719	2,526	49,232	77%
2015	29	1,181,139	38,901	8,181	58,150	420	3,098	50,805	85%

Table 1. Raw IPv6 statistics and deployment over 7 years of data (January, 1st of each year), where "VPs" gives the number of probing monitors, "Traces" the amount of `traceroute` performed, "prefixes" the number of probed prefixes, "ASes" the amount of different ASes in the dataset, "Addresses" the number of pure IPv6 addresses, IPv4-mapped IPv6 addresses and addresses involved in MPLS IPv6 tunnels, and "Tunnels" provides the number of unique MPLS tunnels encountered (note that "Complete Tunnels" refer to tunnels where all LSRs responded to `traceroute` probes).

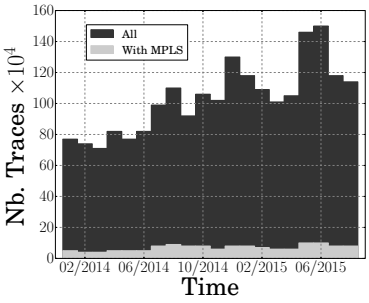

Fig. 2. Raw number of IPv6 traces traversing at least one MPLS tunnel.

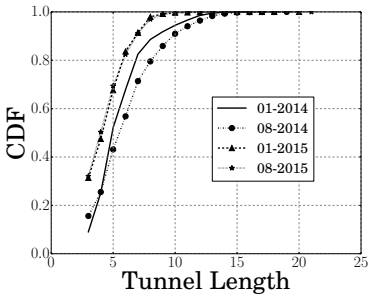

Fig. 3. IPv6 MPLS tunnels length distribution

be replaced by a *. If the router has no IPv6 connectivity but is IPv6-aware, it may send an ICMPv6 message, using a so-called IPv4-mapped IPv6 address [20] as source address. The error message is then propagated towards the Egress LER using MPLS, and then propagated through IPv6 routing.

3 Evaluation

3.1 Dataset

For evaluating the deployment of MPLS under IPv6, we use the IPv6 Archipelago dataset [21]. The data is collected by performing ICMP-based ParisTraceroute measurements [22], using `scamper` [23]. Each vantage point probes all announced IPv6 prefixes (/48 or shorter) once every 48 hours by targeting a single random destination in each prefix. Some vantage points might, in addition, probe the first address (i.e., ::1) in a prefix.

Table 1 provides raw statistics about the IPv6 dataset. We collected the probing campaign made every January 1st since 2009. From this dataset, we extracted the various traces, explicit MPLS tunnels, and performed an IP2AS mapping using Team Cymru.[3] As already stated by others [2, 1], we observe a slow deployment of IPv6 between 2009 and 2013, compensated by a fast increase between 2014 and 2015. MPLS deployment in IPv6 follows that tendency, the peak of MPLS tunnels being reached in 2014 and 2015. In the following subsection, we will focus on data collected between January, 1st 2014 and August, 1st 2015. For that period of time, we take into account the first measurement snapshot of each month, leading to 20 measurement cycles.

Fig. 2 and Fig. 3 provide basic statistics about MPLS deployment in IPv6. In particular, Fig. 2 gives the raw number of `traceroute` (between January, 1st 2014 and August, 1st 2015) that traverses at least one MPLS tunnel. If the quantity of traces increases over time, on the contrary, the amount of traces involved in an MPLS tunnel remains quite stable. Compared to IPv4 [4, 5, 7], `traceroute` are traversing much less MPLS tunnels: on the order of 7-8% in IPv6 against (at least) 40% for IPv4. Note that the drop observed, in terms of number of traces seen, in early 2015, is due to less active vantage points.

Fig. 3 gives the tunnel length distribution for four measurement snapshots, including Ingress and Egress LER in the length distribution. This means that a length of 3 corresponds to a tunnel with a single LSR. We observe that the tunnel length oscillates between 3 and 21. More interestingly, the tunnel length seems to decrease over time, i.e., tunnels observed in 2015 are shorter than tunnels in 2014. This is due to the fact that MPLS tunnels of AS174 (Cogent) disappear from the dataset around October 2014, and Cogent made use of long tunnels.[4] While encountering a few longer tunnels, MPLS IPv6 tunnels length distribution follows observations made in IPv4 [4, 5].

3.2 Label Stack Size Distribution

In this section, we study the characteristics of IPv6 MPLS tunnels compared to IPv4 ones. First, we are interested in the typical LSE size used by both data planes (i.e., the number of MPLS labels contained at each single LSR). The methodology we follow is quite simple: for each LSR of each tunnel, we count the number of labels contained in the stack and, on this basis, we map each tunnel to the maximum number of labels revealed by each of its LSRs. For short MPLS tunnels, it allows for mapping them to their most likely usage. For instance, and for IPv4 data plane, a short tunnel made of three LSRs such that we find the sequence 1,2,1 (in term of LSE sizes) we claim that such a tunnel is likely to be used for VPN purposes so that we retain the maximum value of two to map it to a 2-label LSP. Note that, in such a case, the bottom label is constant (i.e., the same from end-to-end) in order to denote the outgoing VRF (Virtual Routing and Forwarding) to use at the Egress LER.

[3] `http://www.team-cymru.org/`.

[4] The impact of AS174 in IPv6 has already been discussed in the past [24, 25].

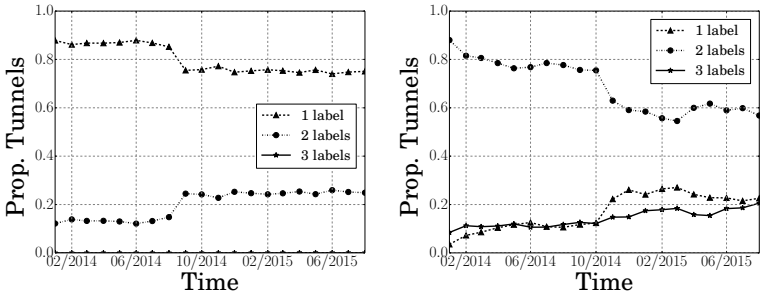

(a) IPv4 (b) IPv6

Fig. 4. LSE stack size distribution over time.

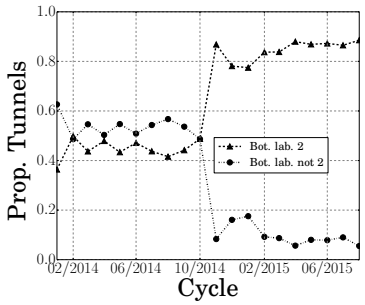

Fig. 5. Distribution of the value in the bottom stack LSE in IPv6.

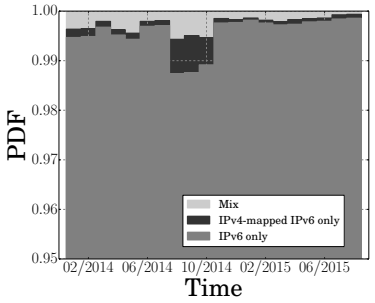

Fig. 6. 6PE core architecture.

Fig. 4 shows the LSE size distribution over time, between January, 1st, 2014 and August, 1st, 2015. Globally speaking, we observe a different behavior between IPv4 (Fig. 4(a)) and IPv6 (Fig. 4(b)). Indeed, under IPv4, the majority of tunnels (around 80%) exhibits a single LSE (this results is aligned with Sommers et al. observations [4]) while, in IPv6, the majority of tunnels (around 80% also during the first half of the considered period) exhibits at least two labels.[5] This result may appear really surprising since there is no obvious reason that justifies a more extensive use of VPN in IPv6 than in IPv4.

Fig. 5 deeper investigates the LSE typical content in IPv6. In particular, it looks at the label value at the bottom of the stack. As explained in Sec. 1, if the value is 2, it suggests a usage of 6PE where core LSRs are dual-stack capable.

Fig. 5 clearly depicts a shift around October 2014. Before that date, tunnels with a label stack are observed almost as often with a bottom label 2 as with another bottom label. After October 2014, things are crystal clear: the majority of tunnels (more than 80%) having a LSE stack use a bottom label of 2, meaning the usage of this type of 6PE is prevalent.

[5] The drop around October 2014 in IPv6 is due to the drastic decrease of MPLS usage by Cogent in the dataset. We show in details in Sec. 3.3 that Cogent has been a heavy user of LSE stacks but then got rid of MPLS.

For tunnels with bottom label 2, we remove this bottom label, and compare the series of remaining labels with series of labels from tunnels found in IPv4 MPLS traces (in the same period).[6] We find out that a match is present in more than 40% of the cases meaning that the same IPv4 LSP is used for IPv6 traffic reinforcing so our assumption about the 6PE usage.

The radical behavior change depicted in Fig. 5 is very surprising at the first glance. To explain it, we investigate the different ASes we encountered in the dataset around this date. Before October 2014, around 50% of the tunnels belong to AS174 (Cogent). In November 2014, this AS disappears from the MPLS dataset while it remains visible through numerous non MPLS IPs. Almost all tunnels belonging to this Tier-1 network have a 2-label stack, but never use the bottom label 2. This is the reason why we have such a behavior modification in Fig. 5. The usage of label stack for Cogent is investigated in details in Sec. 3.3 since it is almost only specific to this AS.

Fig. 6 looks at the architecture of the network core in case of dual-stack 6PE usage (i.e., bottom label 2). We observe a tiny proportion of 6PE tunnels that map IPv4 addresses into IPv6 ones (black region in Fig. 6). We understand this as a case where core LSRs are IPv6 aware (i.e., they are dual-stack) but do not have public IPv6 addresses. In order to be able to reply to probes (i.e., generating ICMP `time-exceeded` messages), they map their IPv4 address in a "fake" IPv6 one.

Most dual-stack 6PE tunnels we observed in the dataset have an IPv6 core (LSRs are dual-stack and have public IPv6 addresses). However, this 6PE usage corresponds to the case where LSPs are deployed with LDPv4. That is, the same LSPv4 generally built with LDP and attached to a given IPv4 loopback destination on the Egress LER is used for both IPv4/v6 traffics. The bottom label 2 indicates to the Egress LER that the traffic is made of IPv6 packets rather than IPv4 ones (where the LSP is made of the same series of top labels without the bottom label). However, note that in practice and at the origin, 6PE has been used to ensure connectivity between IPv6 islands with a tunnel having a pure IPv4 core. In this case, LSRs are not IPv6 aware (no dual-stack and no IPv6 addresses). In such an architecture, IPv4 LSRs will not respond to IPv6 probes of `traceroute`, and the traces in our dataset are incomplete (several `*` appear between 6PEs). Unfortunately, we are not able to differentiate such a behavior from IPv6 nodes that simply not respond to the probes. The proportion of this type of 6PE tunnels is therefore underestimated in this paper.

3.3 The Cogent Case

The Cogent case is particularly interesting and quite intriguing. It has both a very prominent position in the dataset (Cogent is one of the largest Tier-1, in particular it has the second highest AS rank according to CAIDA[7]) and, most of all, an MPLS IPv6 behavior completely different from other ASes we observed.

[6] These IPv4 MPLS traces were also downloaded from the Archipelago dataset.
[7] See `http://as-rank.caida.org/`

This can be seen in Fig. 5 where around October 2014 Cogent more or less disappears from CAIDA MPLS traces and at the same time the proportion of stacks with bottom label 2 rises sharply.

First, the fact that MPLS traces are almost absent from Cogent after this date may either be due to a change in the configuration of its routers or, more simply, that the operator gets rid of MPLS. We conduct some tests to understand whether Cogent removes `ttl-propagate` at Ingress LERs to make MPLS tunnels invisible or not. This type of phenomenon has already been observed for IPv4 MPLS (look at Vanaubel et al. [7] and the specific study on Level3). For this purpose, we pick a subset of MPLS traces obtained before October 2014 and try to find similar pure IP traces after this date (i.e., we check whether the same sequence of IP addresses between Ingress and Egress LER exists before and after that date or if the two edge routers seem directly connected after). As a result, we find equivalent traces before and after, the only difference being that MPLS labels disappear after October 2014. We can conclude that Cogent just gets rid of MPLS (as they did in IPv4 two years before). To verify this conclusion, we contacted a Cogent network administrator who confirmed this first result.

The second, and most interesting fact is that, although most of its LSPs have a stack size greater or equal than 2, they never use a bottom label of 2 (the default value for 6PE), on the contrary to the dominant usage in other ASes (see Sec. 3.2). Note that RFC4798 [18] does not mandate the use of label 2 as bottom label but that BGP at the Egress router associates a label to each IPv6 prefix and announces it to its iBGP peers. Therefore, a 6PE implementation could choose any other arbitrary label for 6PE, or choose a different label for each prefix or set of prefixes.

After the analysis of Cogent stacks, it appears that the bottom label is not fixed (Cogent does not simply use another arbitrary value than 2) but varies greatly. In fact numerous different bottom labels can be found on LSPs connecting the same (Ingress, Egress) pair. For instance, we find one case where 38 distinct bottom labels are in use for a given pair. In theory, this could be 38 distinct VPNs or, more probably, the Egress could be using a distinct bottom label for each (group of) IPv6 prefix. Hopefully, our Cogent contact helped us to eliminate the VPN case (indeed, considering only the measurements perspective, nothing distinguishes VPN from 6PE, the general principle of using a bottom label being the same): Cogent simply did not use this technology but only 6PE before shutting down MPLS for IPv6 in October 2014.

One purpose of distinct bottom labels may be load sharing: in a network using Equal Cost Multipath (ECMP), packets arriving at a router with two equal cost routes for the destination are distributed along these routes according to a packet header hash. In a network using MPLS and ECMP, LSPs constructed by LDP signaling may make use of multiple paths, building several LSPs between the same pair of LSRs. In the case of Cogent, it is apparent that ECMP is in use in the core network. For example in the case of the Egress router having 38 distinct bottom labels, after removing the bottom label, there still exists 8 distinct LSPs between this pair of routers (considering IP addresses and top labels).

	09/2014	10/2014	11/2014
Mono-LSP	23.1%	30.8%	0%
Multi-FEC	3.4%	2.7%	0%
Mono-FEC	58.6%	52.3%	0%
Unclassified	14.9%	14.2%	0%

Table 2. LPR [7] applied to some Cogent IPv6 2014 data.

For IPv4 packets, the hash function considers at least the IP addresses in order to guarantee the same route for all packets of the same flow (avoiding so ordering issues with TCP). The same can be done with IPv6 packets, but it is more costly due to the IPv6 addresses length. Moreover in the case of 6PE, routers in the core may be totally IPv6 ignorant. In this case using the bottom label to split the load makes sense (this usage is for example mentioned in Cisco documentation [26]). Note that the conjunction of many routes (ECMP), hence top labels and many distinct 6PE bottom labels result in a large number of distinct LSPs when taking into account the full label stack. This partially explains why Cogent is so prevalent in terms of unique IPv6 MPLS tunnels in the dataset we consider besides its shere size. Several mechanisms have been proposed to allow MPLS networks to benefit from the use of multiple paths, such as Kompella et al. [27]. There have been also proposals to allow RSVP-TE to make a direct use of multiple ECMP paths [28].

To investigate further and retrieve the root cause of this variety of label stacks, we apply the Label Pattern Recognition (LPR) algorithm [7] on top labels of the Cogent IPv6 MPLS traces to quantify the usage of LDP (for IGP-BGP scalability purposes – *Mono-FEC* in Table 2) and/or RSVP-TE (for traffic engineering purposes – *Multi-FEC* in Table 2). To distinguish LSPs built through LDP and RSVP-TE, LPR analyses LSPs going through the same Ingress-Egress pair. If two LSPs have been built through LDP, the incoming top labels should be identical at converging routers. On the contrary, the incoming top labels should be different if these LSPs have been built through RSVP-TE. There is also the possibility that both protocols are used, building different LSPs according to the intended service. Our analysis (already apparent in the case of the Egress with 38 distinct LSPs), shows that the top-label is mostly generated by LDP (Mono-FEC line in Table 2). Therefore our interpretation is that the bottom-label is assigned by the Egress-router on a per IPv6 prefix basis using a variant of 6PE, in order to make a more efficient use of ECMP, while the top-label (i.e. the LSP itself) is built using LDP for IPv4.

4 Conclusion

The recent years have seen an increasing deployment and usage of IPv6. With the recent IPv4 depletion, this increase is going faster and we expect to see more and more IPv6 networks in a near future. In this paper, we focused on a specific aspect of the IPv6 deployment related to MPLS: how is MPLS deployed and

used under IPv6? Is its usage strongly different from the one in IPv4? Based on `traceroute` collected by CAIDA, we tried to answer these questions.

Our first observations pointed out that the MPLS technical usage seems to strongly differ between IPv4 and IPv6. In particular, in the way label stacks are used, we discovered that under IPv4, stacks of more than one label are not that frequent while they are the norm under IPv6. However, we showed that this difference is not due to an increase in VPN BGP MPLS usage. Indeed, we explained that IPv6 MPLS mostly uses 6PE tunnels that are built using an IPv4 signaling protocol (in particular LDP for IPv4). This allows one to deploy MPLS for IPv6 across a network where some routers are not dual-stack, or where LDP is not available for IPv6 (the IPv6 version was only recently released). Therefore this can be seen as a transition mechanism, and it will be interesting to see the evolution of this usage as more and more networks become fully IPv6 compliant. The special case of the Cogent network also brought some light on the use of ECMP multipath in conjunction with MPLS. We argued that this network uses a specific form of 6PE to ease the way that IPv6 routers select their outgoing interfaces.

References

1. Dhamdhere, A., Luckie, M., Huffaker, B., claffy, k., ELmokashfi, A., Aben, E.: Measuring the deployment of IPv6: Topology, routing, and performance. In: Proc. ACM Internet Measurement Conference (IMC). (November 2012)
2. Czyz, J., Allman, M., Zhang, J., Iekel-Johnson, S., Osterweil, E., Bailey, M.: Measuring IPv6 adoption. In: Proc. ACM SIGCOMM. (August 2014)
3. American Registry for Internet Numbers (ARIN): IPv4 depletion (September 2015) See `https://www.arin.net/resources/request/ipv4_countdown.html`.
4. Sommers, J., Eriksson, B., Barford, P.: On the prevalence and characteristics of MPLS deployments in the open Internet. In: Proc. ACM Internet Measurement Conference (IMC). (November 2011)
5. Donnet, B., Luckie, M., Mérindol, P., Pansiot, J.J.: Revealing MPLS tunnels obscured by traceroute. ACM SIGCOMM Computer Communication Review **42**(2) (April 2012) 87–93
6. Pathak, A., Zhang, M., Hu, Y.C., Mahajan, R., Maltz, D.: Latency inflation with MPLS-based traffic engineering. In: Proc. ACM Internet Measurement Conference (IMC). (November 2011)
7. Vanaubel, Y., Mérindol, P., Pansiot, J.J., Donnet, B.: MPLS under the microscope: Revealing actual transit path diversity. In: Proc. ACM Internet Measurement Conference (IMC). (October 2015)
8. Rosen, E., Visanathan, A., Callon, R.: Multiprotocol label switching architecture. RFC 3031, Internet Engineering Task Force (January 2001)
9. Andersson, L., Asati, R.: Multiprocotol label switching (MPLS) label stack entry: EXP field renamed to traffic class field. RFC 5462, Internet Engineering Task Force (February 2009)
10. Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., Conta, A.: MPLS label stack encoding. RFC 3032, Internet Engineering Task Force (January 2001)

11. Agarwal, P., Akyol, B.: Time-to-live (TTL) processing in multiprotocol label switching (MPLS) networks. RFC 3443, Internet Engineering Task Force (January 2003)
12. Andersson, L., Minei, I., Thomas, T.: LDP specifications. RFC 5036, Internet Engineering Task Force (October 2007)
13. Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., Swallow, G.: RSVP-TE: Extensions to RSVP for LSP tunnels. RFC 3209, Internet Engineering Task Force (December 2001)
14. Muthukrishnan, K., Malis, A.: A core MPLS IP VPN architecture. RFC 2917, Internet Engineering Task Force (September 2000)
15. George, W., Pignataro, C.: Gap analysis for operating IPv6-only MPLS networks. RFC 7439, Internet Engineering Task Force (January 2015)
16. Asati, R., Pignataro, C., Raza, K., Manral, V., Papneja, R.: Updates to LDP for IPv6. RFC 7552, Internet Engineering Task Force (June 2015)
17. De Clercq, J., Ooms, D., Carugi, M., Le Faucheur, F.: BGP-MPLS IP virtual private network (VPN) extension for IPv6 VPN. RFC 4659, Internet Engineering Task Force (September 2006)
18. De Clercq, J., Ooms, D., Prevost, S., Le Faucheur, F.: Connecting IPv6 islands over IPv4 MPLS using IPv6 provider edge routers (6PE). RFC 4798, Internet Engineering Task Force (February 2007)
19. Bonica, R., Gan, D., Tappan, D., Pignataro, C.: ICMP extensions for multiprotocol label switching. RFC 4950, Internet Engineering Task Force (August 2007)
20. Hinden, R., Deering, S.: IP version 6 addressing architecture. RFC 4291, Internet Engineering Task Force (February 2006)
21. CAIDA: The CAIDA UCSD IPv6 topology dataset (September 2015) See `http://www.caida.org/data/active/ipv6_allpref_topology_dataset.xml`.
22. Augustin, B., Cuvellier, X., Orgogozo, B., Viger, F., Friedman, T., Latapy, M., Magnien, C., Teixeira, R.: Avoiding traceroute anomalies with Paris traceroute. In: Proc. ACM Internet Measurement Conference (IMC). (October 2006)
23. Luckie, M.: Scamper: a scalable and extensible packet prober for active measurement of the Internet. In: Proc. ACM Internet Measurement Conference. (November 2010)
24. Giotsas, V., Luckie, M., Huffaker, B., claffy, k.: IPv6 AS relationships, clique, and congruence. In: Proc. Passive and Active Measurement Conference (PAM). (March 2015)
25. Leber, M.: IPv6 Internet broken, Cogent/Telia/Hurricane not peering (October 2009) Nanog Mailing-list. See `http://mailman.nanog.org/pipermail/nanog/2009-October/014017.html`.
26. Cisco: Cisco IOS IPv6 provider edge router (6PE) over MPLS (October 2015) See `http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_data_sheet09186a008052edd3.html#wp39913`.
27. Kompella, K., Drake, J., Amante, S., Henderickx, W., Yong, L.: The use of entropy labels in MPLS forwarding. RFC 6790, Internet Engineering Task Force (November 2012)
28. Kompella, K., Hellers, M., Singh, R.: Multi-path label switched paths signaled using RSVP-TE. Internet Draft (Work in Progress) draft-kompella-mpls-rsvp-ecmp-06, Internet Engineering Task Force (March 2015)