Evaluating the Benefits of the Locator/Identifier Separation*

Bruno Quoitin, Luigi Iannone, Cédric de Launois, Olivier Bonaventure
Université Catholique de Louvain (UCL)
Belgium
{firstname.lastname}@uclouvain.be

ABSTRACT

Since recent years, it has been recognized that the existing routing architecture of today's Internet is facing scalability problems. Single numbering space, multi-homing, and traffic engineering, are making routing tables of the default free zone to grow very rapidly. Recently, in order to solve this issue, it has been proposed to review the Internet addressing architecture by separating the end-systems identifiers' space and the routing locators' space.

In this paper we review the most recent Locator/ID separation proposal and explore the benefits that such an architecture may bring. In particular, we evaluate the improvements that can be achieved in terms of routing tables' size reduction and traffic engineering.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Network communications; C.2.6 [**Internetworking**]: Routers; C.4 [**Performance of Systems**]: Design studies

General Terms

Management, Measurement, Performance, Design, Experimentation, Standardization.

Keywords

Locator ID separation, Traffic Engineering, Routing, Addressing.

1. INTRODUCTION

Last years have witnessed an increasing concern about the current IP routing and addressing architecture, perceiving that the use of a single numbering space, namely the *IP addressing space* ([6, 24, 14]), for both host transport session identification and network routing creates scaling issues. Multi-homing, Traffic Engineering (TE), and suboptimal address allocation are making the Forwarding Information Base (FIB) of the Default Free Zone (DFZ) growing at a greater than linear rate [14]. BGP's Tables have already reached 200,000 entries. Furthermore, the peculiarity of BGP, advertising only one route for each prefix, introduces strong limitations in performing inter-domain TE.

Recent discussions in the IETF and IRTF [17] suggest that scaling benefits could be realized by separating the current IP address space into separate spaces for end-systems identifiers and routing locators. Among these benefits, we can mention the following.

- *Reduction of routing table size in the DFZ.* Differently from today's current practice, where addresses are more and more assigned in a provider independent way, the use of a separate numbering space for routing locators will allow to assign them in a topologically driven manner. In turn this would allow a high level of aggregation, reducing the number of globally announced prefixes.
- *Improved Traffic Engineering capabilities.* Today, TE is often achieved by de-aggregating IP prefixes. By separating ID and locators it is possible to perform both inbound and outbound flexible TE, setting tunnels between locators based on several different metrics or policies. Furthermore, traffic between end-systems and routing locators can be redistributed by taking advantage of the identifier-to-routing-locator mapping function.

Despite some divergences, the community seems to agree that this Locator/ID separation is a basic component of the future Internet architecture ([19, 18, 20, 28, 3]). Recently, a protocol called LISP (which stands for Locator/ID Separation Protocol [10]) has been proposed to support the incremental deployment of this separation.

In this paper we explore the above mentioned benefits. We first describe how this separation between locators and identifiers can be achieved using LISP. Then, based on both real measurements and simulations, we evaluate the benefits that such an approach may enable. Starting from a realistic Internet topology, we also explore what would be the impact on the routing tables when using only aggregatable addresses for locators. Furthermore, we evaluate the path's diversity inherently present in the Internet, i.e., the number of alternative Internet's routes to reach the same destination domain. Currently, these routes are not exploited, due to the BGP's characteristic of advertising only one (best) route. Given this path diversity, we show how it is possible to take advantage of it to reduce end-to-end latency.

Remark that, the separation between locators and ID ease the migration of stub networks from one provider to another, without

---

*This work was partially supported by the European Commission within the IST-027609 AGAVE project (`www.ist-agave.org`) and by the Walloon Region within the TOTEM project (`totem.info.ucl.ac.be`)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiArch'07, August 27–31, 2007, Kyoto, Japan.

Copyright 2007 ACM 978-1-59593-784-8/07/0008 ...$5.00.

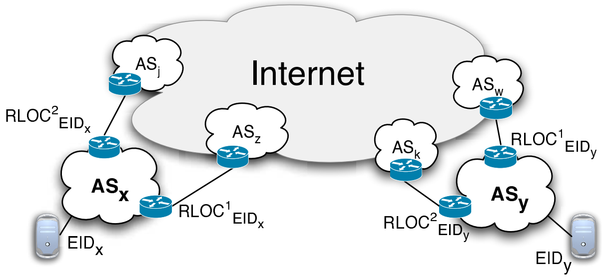

Figure 1: Position of EIDs and RLOCs in the global Internet.

the need of renumbering. If a dynamic binding is used between stub networks and locators, locator/identifier separation can be used to manage network mobility, like in NEMO [9].

The paper is organized as follows. In section 2 we give a brief overview of the LISP protocol and how the separation between ID and Locators can be achieved. In section 3 we evaluate the reduction of the FIBs' size that can be obtained in the DFZ. In section 4 we quantify the unexploited path diversity present in today's Internet, while in section 5 we show an example of how to take advantage of this characteristic.

2. LOCATOR/ID SEPARATION WITH LISP

The Locator/ID Separation Protocol (LISP) [10] is a simple IP-over-IP tunneling protocol aiming at giving a network layer support to routing locators and end-host identifiers separation. A key contribution of LISP is that it can be incrementally deployed. Other works present in the literature have the same or a similar target (e.g. [13, 26, 15, 21]), however, they mainly have a disruptive impact on the current architecture and/or need heavy changes in the protocol stack of end-systems.

On the contrary, a main objective of LISP is to provide Locator/ID separation without the need of modifying in any way the current protocol stack of today's end-systems. This is one of the major requirements that have been pointed out in recent discussions in the research community. End-systems will still send and receive packets using IP addresses, which in the LISP terminology are called Endpoint Identifiers (EIDs). LISP is defined in four different variants, depending on the *routability* of EIDs. In the first two variants of LISP (called LISP 1 and 1.5), EIDs will still be IP routable addresses, in order to facilitate incremental deployment. Nevertheless, in the future (i.e., variants 2 and 3 of LISP), EIDs will be only locally routable IP addresses (i.e., that are routable only in the local AS, similar to IPv6 Site Local Addresses). In order for EIDs to send/receive packets from outside the local AS, they are associated to one or more *Tunnel Routers*, whose IP address is called Routing Locator (RLOC) in the LISP terminology. RLOCs are, and will also remain in the future, globally routable IP addresses associated to the Tunnel Routers through which EIDs can be reached.

In order to explain how the separation of EIDs and RLOCs works, in particular when using LISP, let us take as example the topology depicted in figure 1. The end-host $\mathsf{EID_x}$ is reachable through two border routers, meaning that it can be associated to two locators: $\mathsf{RLOC^1_{EID_x}}$ and $\mathsf{RLOC^2_{EID_x}}$.[1] Similarly, $\mathsf{EID_y}$ has two locators: $\mathsf{RLOC^1_{EID_y}}$ and $\mathsf{RLOC^2_{EID_y}}$.

During end-to-end packet exchange between two Internet hosts, an Ingress Tunnel Router (ITR) prepends a new LISP header to each packet and the Egress Tunnel Router (ETR) strips this new header before delivering the packet to the final destination. Remark that the LISP header is a normal IP header with the only peculiarity of using locators as source and destination addresses. For instance, in the case of figure 1, assuming that $\mathsf{EID_x}$ wants to open a connection to $\mathsf{EID_y}$ the following steps are performed.

1. $\mathsf{EID_x}$ issues a first IP packet, using its ID ($\mathsf{EID_x}$) as Source Address (SA) and using $\mathsf{EID_y}$ as Destination Address (DA). This packet is routed inside $\mathsf{AS_x}$ in the usual way, in order to be delivered to one of $\mathsf{EID_x}$'s locators.

2. The ITR ($\mathsf{RLOC^1_{EID_x}}$) receives the packet. Remark that this choice is done in practice by intra-domain TE policies coherently to the local EID-to-RLOC mapping. These policies can vary from AS to AS. Nonetheless, the EID is still reachable from outside through all of its RLOCs.

3. $\mathsf{RLOC^1_{EID_x}}$ performs EID-to-RLOC lookup to determine the locator of $\mathsf{EID_y}$, and, thus, the corresponding routing path through which the packet will be forwarded. Assuming that this operation returns $\mathsf{RLOC^2_{EID_y}}$, the EID-to-RLOC association is kept in a dedicated cache.

4. A LISP header is prepended to the original IP packet, having $\mathsf{RLOC^1_{EID_x}}$ as SA and $\mathsf{RLOC^2_{EID_y}}$ as DA.[2] The packet is then routed at IP level inside the Internet.

5. Once the LISP packet reaches $\mathsf{RLOC^2_{EID_y}}$ the LISP header is stripped and the packet is forwarded inside $\mathsf{AS_y}$ as usual, in order to be delivered to $\mathsf{EID_y}$.

Note that, by comparing the stripped LISP header with the inner IP header, $\mathsf{RLOC^2_{EID_y}}$ is able to retrieve the EID-to-RLOC association of the sender ($\mathsf{EID_x}$), which can be stored in the local cache. No reverse mapping lookup is needed. After the first packet has gone through, the caches on both endpoints of the LISP tunnel have the appropriate information to correctly forward all the subsequent packets.

As of this writing the EID-to-RLOC mapping function, which is a main component of the Locator/ID separation paradigm, is still object of discussion [5]. A non-exhaustive list of proposals include relying on new ICMP control messages to discover the set of RLOCs of a given EID (variant 1 of LISP [10]), relying on the DNS service (variant 2 of LISP [11]), on overlay networks [25, 27] and on BGP [22]. In this paper, we do not tackle this issue and do not propose a new EID-to-RLOC mapping mechanism. Instead, we discuss and evaluate the benefits that the Locator/ID separation can bring in terms of routing scalability and traffic engineering opportunities.

3. SHRINKING THE FIB

The Locator/ID separation provided by LISP allows reducing the size of FIBs. More compact FIBs implies lower memory requirements for routers, possibly faster lookups and faster forwarding table updates. The FIB size reduction is possible since locators are now independent of identifiers. They can therefore be allocated in

---

[1] In the sake of simplicity, we use the same acronyms to indicate both the name of the system and its IP address, i.e., both $\mathsf{EID_x}$ and $\mathsf{RLOC^2_{EID_y}}$ indicates at the same time a name and an IP address.

[2] In LISP version 1 and 1.5, due to incremental deployment purposes, actually the DA is set to $\mathsf{EID_y}$, however, this will not be the case in LISP 2 and 3. For more details please refer to [10].

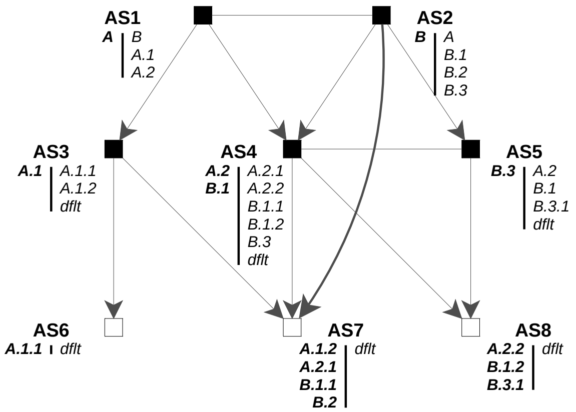

Figure 2: Scenario S1 deployed on a small, example topology.

a more aggregatable way than with today's IPv4 prefixes without impeding on the customer's freedom to change their providers. In the following we evaluate this reduction by analyzing two different scenarios for allocating the global prefixes in the Internet in a more aggregatable way.

The first scenario, called **S1**, assumes that only Tier-1 domains have been allocated a globally advertised prefix. Each Tier-1 delegates a non-overlapping fraction of its own prefix to each of its customers. The non-Tier-1s-transit domains subsequently allocate to each of their customers a fraction of each of their own prefixes. In **S1**, providers can advertise a single default route to their customers. On shared-cost peerings, only routes towards prefixes owned by the peers are exchanged, since all the customer routes are aggregated. This scenario is illustrated on a simple example in figure 2. In the example, edges with an arrow depict provider-customer relationships while edges without arrow depict peer-to-peer relationships. Besides each AS in figure 2, we have shown the received prefixes (in bold) and the FIB entries. For instance, AS7 has AS2, AS3, and AS4 as providers and has received prefixes A.1.2, A.2.1, and B.2. Moreover, AS7 needs a single FIB entry that corresponds to its default route through one of the providers.

Since this approach of scenario **S1** can lead some domains to be allocated a large amount of prefixes, we also investigate variants limiting to 2 or 5 the number N of prefixes delegated to the customers. As we will show later, even limiting the number of prefixes delegated to the customers, full connectivity is still guaranteed while greatly reducing the FIBs' size.

In the second scenario, called **S2**, the hierarchical allocation of prefixes is less stringent. We assume that all the transit domains (Tier-1s and non-Tier-1s) are allocated a globally-advertised prefix. These prefixes are fractioned and assigned to the customers. In this scenario, there are a larger number of globally-advertised prefixes but customers are allocated a single prefix per provider they connect to. A simple example of Scenario **S2** is illustrated in figure 3. In this case, all the transit ASes (AS1 to AS5) have received independent prefixes (A to E, respectively). These prefixes are advertised globally. Stub ASes only receive a fraction of their providers' prefixes.

Name	ASes	T1s	Transits	Stubs	Depth
Large	14965	2	2707	11986	9
Small	11923	50	0	11873	1

Table 1: Parameters of the topologies

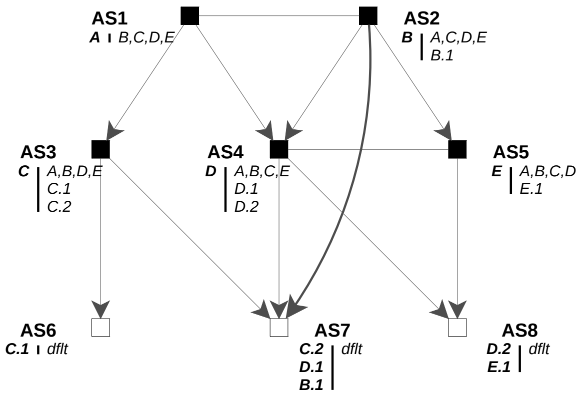

Figure 3: Scenario S2 deployed on a small, example topology.

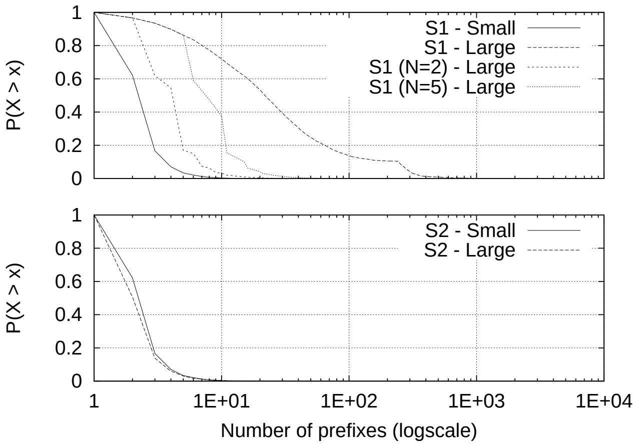

Figure 4: Number of per domain assigned prefixes.

In order to compare these scenarios, we simulated them on two large topologies. We used GHITLE [8] to generate Internet-like AS-level topologies. GHITLE relies on a preferential attachment algorithm and assigns each edge in the topology a business relationship (*provider-to-customer* or *shared-cost*). Table 1 shows the number of ASes and the breakdown between Tier-1s/Transits/Stubs in each topology. The **Large** topology is an Internet-like topology with a small number of Tier-1s, a lot of stubs and a significant number of transit domains in-between. The **Small** topology has only two levels: Tier-1s and customers (stubs).

We show in figure 4 the number of prefixes that each domain is being allocated in both scenarios. In **S1**, a large fraction of the domains (around 15%) receive more than 100 different prefixes, while in **S2** the majority of the domains receive a single prefix and 99.9% receive less than 10 prefixes. In **S1**, the domains with a large number of prefixes are typically well-connected regional access networks, i.e. located at the bottom of the hierarchy and with a large number of providers. For scenario **S1**, we also plotted the results when the number N of prefixes that an AS delegates to its customers is limited to 2 and 5. In this case, the number of prefixes that any domain can receive is bounded by N times the number of its providers.

In figure 5, we show the number of FIB entries for external destinations in each domain. We observe that in both scenarios, more

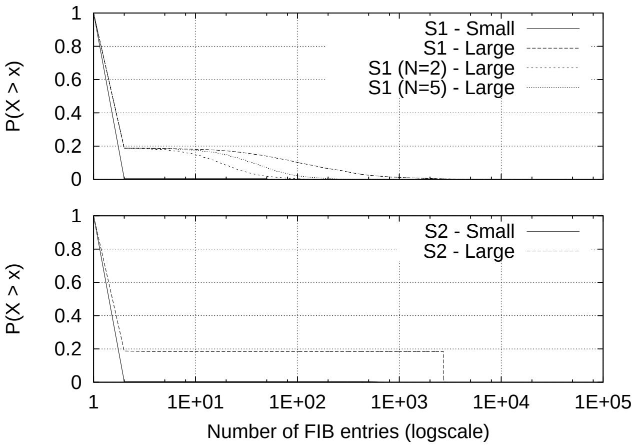

Figure 5: Number of per domain installed FIB entries.

than 80% of the domains need only a single FIB entry. These are the stubs that only need a default exit route. The main difference between **S1** and **S2** concerns the transit domains. In **S2**, all the transit domains need a different FIB entry for all the other transit destinations and a FIB entry for each of their non-transit customer. Roughly, this means that there are at least 2708 entries in the FIB of transit domains, one per each other transit domain. On the contrary, in **S1**, there are fewer globally-advertised routes and therefore less need for large FIBs. The major contribution to the FIB comes from the customer and shared-cost entries.

The results presented above show that distributing RLOCs in an aggregatable way allows to strongly reduce FIBs' size. Figure 5 in particular shows that the number of entries has an order of magnitude less than the number of ASes, while in the current Internet the number of entries has an order of magnitude larger than the number of ASes. Allocating prefixes in a more aggregatable way also reduces the RIBs size as well as the constant churn of BGP messages since fewer destinations are advertised in the default-free zone.

4. ROUTE DIVERSITY IN THE INTERNET

The support for Locator/IP separation provided by LISP enables the possibility to perform more flexible inter-domain TE. The EID-to-RLOC lookup operation can provide a list of RLOCs from which one can choose in order to optimize some performance metrics.

For instance, in the example of section 2, the lookup operation performed by $\mathsf{RLOC}^1_{\mathsf{EID}_x}$ can return both $\mathsf{RLOC}^1_{\mathsf{EID}_y}$ and $\mathsf{RLOC}^2_{\mathsf{EID}_y}$ locators for EID_y. At this point AS_x can choose one of them based on some optimization criteria (e.g. delay). Remark that in the current Internet there is not such flexibility due to the BGP's characteristic of advertising only one route to each AS.

In order to evaluate the potential benefit of exploiting the multiple routes towards the providers of the destination domain, we performed a simulation based on real BGP routing tables collected by the Route Views Project [16]. The study [21] has been performed on a routing table collected on December 1st, 2004. The routing table contained 5750380 routes received from 34 different peers. In the simulation, we only considered the 32 peers that announced a full routing table, i.e., more than 140000 routes.

Among all the received routes, we identified, based on the AS-paths, 6402 multihomed stubs. These multihomed stubs originated 29575 different prefixes. We then considered all the 496 pairs of RouteViews peers. For each pair of peers, we simulated a dual-homed stub domain connected to the peers. For each simulated stub, we counted the number of different paths learned through BGP towards all the considered destination prefixes, assuming that each provider advertises one prefix per stub. This is similar to the scenario **S2** presented in section 3. Further, we consider that two paths are different if at least the provider in the source AS or the provider in the destination AS is different. Note that if two paths are different, that does not mean that they are completely disjoint [8].

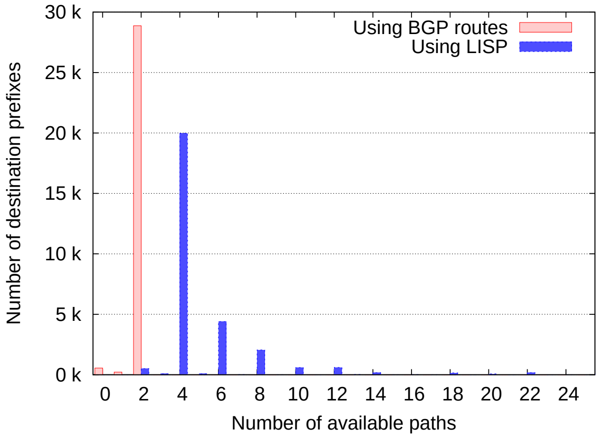

Figure 6: Path diversity when multihoming to RouteViews peers.

The results of our simulations are summarized in figure 6. The figure shows the distribution of the number of different paths available with BGP towards the destination domain vs. the paths towards the destination AS and passing through different ETRs, for all the destination prefixes. On the x-axis, we show the number of different paths available and on the y-axis, the number of prefixes that could be reached with the corresponding number of paths. The number of available paths is an average over the 496 simulated dual-homed stubs. We do not show the variance since it is very low.

When looking at the BGP paths towards the destination AS, the number of distinct paths is comprised between 0 and 2. If there is no path, that means that the destination prefixes cannot be reached. This fortunately occurs for only a small subset of the RouteViews dataset. This is probably due to the filters used by some ISPs. If there is only one path, this means that the destination prefix was not reachable through one of the providers. But most of the time, the destination prefixes were reachable through both providers. The number of available BGP paths cannot be more than 2 since the simulated dual-homed stubs only receive one BGP route for each destination prefix from each provider. The path diversity is thus low with BGP even if there are two different paths most of the time.

If we look at all the routes towards the destination AS and passing through different providers (which can be exploited using LISP), the path diversity increases a lot. Most destination prefixes (67 %) are reachable through at least 4 different paths. There is also a significant number of destination ASes (30 %) that are reachable through more than 4 paths due to some destination stubs being more than dual-homed. The reason for the large majority of the destination prefixes having an even number of different paths is that the source stub is dual-homed. The same study was performed on routing tables collected by the RIPE RCC [23] and similar results were obtained [21]. Previous studies have also shown that a similar behavior is also present in IPv6 topologies [4].

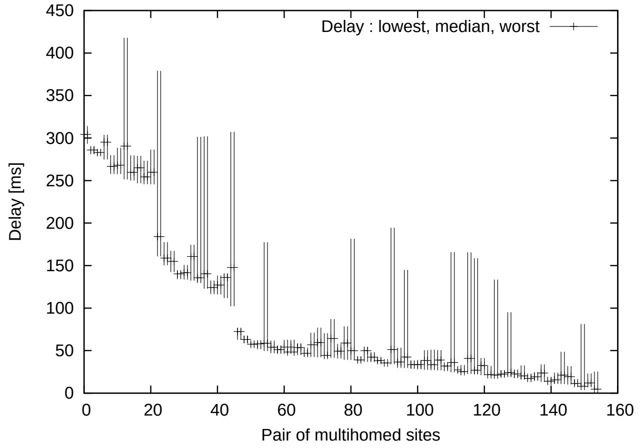

Figure 7: Delays on the available paths between 13 multihomed sites based on RIPE NCC measurements.

5. EXPLOITING ROUTE DIVERSITY

The previous section clearly shows the large amount of path diversity present in the Internet, which is still unexploited. As already stated, Locators/ID separation enables to take advantage of this diversity by smartly selecting RLOCs when several options are available. Here we go deeper in the analysis by evaluating the gain that can be achieved when using delay as the optimization criteria.

We performed a simulation study of the delays along the paths between multihomed sites. The simulation is based on real delay measurements made during May 2004 between 58 active test boxes from the RIPE NCC Test Traffic Measurements Service [12]. The test boxes are scattered over Europe and a few are located in the US, Australia, New Zealand and Japan. Each test box is equipped with a GPS clock so that one-way delays between each pair of boxes can be measured accurately (within 10μs). More than 2000 probes are performed per day and per test box pair. The interval between two consecutive probes is randomized according to a Poisson distribution, as recommended in [2]. All the measures obtained have been then fed to the Vivaldi algorithm [7] in order to obtain stable values of the measured delay, which we used in our simulation.

To simulate the presence of multiple RLOCs, we follow a methodology similar to the one used in [1]. We select a few RIPE nodes in the same metropolitan area, and consider them as the RLOCs of a single virtual multihomed network. This method actually models multihoming where the provider-dependent prefixes advertised by the virtual site are aggregated by its providers, like scenario **S2** in section 3. A total of 13 multihomed sites are emulated by this method: 10 dual-homed sites, 1 three-homed, 1 four-homed, and 1 having 8 providers. One multihomed site is located in the US, one in Japan, and the others in Europe.

Figure 7 shows an analysis of delays between the RLOCs of the 13 multihomed sites. We evaluated the delay for each possible pair (i.e., 78 pairs) of multihomed sites and sorted them in decreasing order of their best delay (cf., x-axis of figure 7). The figure also shows, for each pair, the range of delays on the available paths. In particular, for each pair of multihomed sites, we draw a vertical bar. The upper end of the bar indicates the actual delay of the worst path, while the lower end indicates the actual delay of the best path, i.e. the lowest possible delay. A graduation is added on the bar, indicating the median delay among the paths.

We observe that for many pairs there are large variations in the measured delays, with differences between the best and the worst case larger than 100ms. Due to the performance-blind selection of paths performed by BGP, the worst path could be selected, leading to a delay that can sometimes be tremendously larger than the delay of the best available path.

Locators/ID separation allows choosing among several RLOCs, increasing the freedom of choosing alternative paths such as lower delay paths. Note that delay is not the only possible optimization criteria, but well shows the possible achievable gains.

6. CONCLUSIONS

During the last years, several researchers have proposed mechanisms where locators and identifiers are separated, in contrast to the current Internet architecture. The rationale behind this approach is to overcome the scaling issues that have appeared.

In this paper, we first summarized the behavior of the LISP approach. We have then shown by simulations that such a mechanism helps in significantly reducing the size of the FIB of core Internet routers. This reduction is possible because locators are assigned hierarchically.

Furthermore, the allocation of multiple locators to each stub AS provides additional benefits. We have shown that thanks to these locators, stub ASes can exploit many more paths than when using classical BGP-based multihoming. Our simulations, based on RIPE TTM delay measurements, have also shown that by exploiting more paths stub ASes could obtain paths with a much lower delay.

Acknowledgments

We thank the RIPE NCC for providing the Test Traffic Measurements Service and allowing us to use the collected raw data. We are also grateful to the University of Oregon's Route Views project for providing their BGP routing table dumps.

7. REFERENCES

[1] A. Akella, S. Seshan, and A. Shaikh. Multihoming Performance Benefits: An Experimental Evaluation of Practical Enterprise Strategies. In *Proceedings of USENIX Annual Technical Conference*, Boston, MA, 2004.

[2] G. Almes, S. Kalidindi, and M. Zekauskas. A round-trip delay metric for IPPM. RFC 2681, IETF, September 1999.

[3] O. Bonaventure. Reconsidering the Internet Routing Architecture. *Internet Draft*, March 2007. `http://www.ietf.org/internet-drafts/draft-bonaventure-irtf-rrg-rira-00.txt`.

[4] B. Q. C. de Launois and O. Bonaventure. *Leveraging Network Performances with IPv6 Multihoming and Multiple Provider-Dependent Aggregatable Prefixes*, volume 50. Elsevier Computer Networks, June 2006.

[5] B. Carpenter. General Identifier-Locator Mapping Considerations. *Internet Draft*, June 2007. Available online at: `http://www.ietf.org/internet-drafts/draft-carpenter-idloc-map-cons-01.txt`.

[6] J. Chiappa. Endpoints and enpoint names: A proposed enhancement to the internet architecture. *Internet Draft*, 1999. `http://www.chiappa.net/~jnc/tech/endpoints.txt`.

[7] F. Dabek, F. Kaashoek, and R. Morris. Vivaldi: A decentralized network coordinate system. In *Proceedings of ACM SIGCOMM'04*, Portland, Oregon, USA, August 2004.

[8] C. de Launois. *Unleashing Traffic Engineering for IPv6 Multihomed Sites*. Univeristé Catholique de Louvain, September 2005. PhD Thesis.

[9] V. Devarapalli, R. Wakikawa, A. Petrescu, and P. Thubert. Network mobility (NEMO) basic support protocol. *RFC 3963*, January 2005.

[10] D. Farinacci, V. Fuller, and D. Oran. Locator/id separation protocol (LISP). *Internet Draft*, January 2007. Available online at: `www.ietf.org/internet-drafts/draft-farinacci-lisp-00.txt`.

[11] D. Farinacci, D. Oran, V. Fuller, and J. Schiller. Locator/ID Separation Protocol (LISP2) DNS based version. *Slides-set*, November 2006. Available online at: `http://www.dinof.net/~dino/ietf/lisp2.ppt`.

[12] F.Georgatos et al. Providing Active Measurements as a Regular Service for ISP's. In *Proceedings of PAM'01*, Amsterdam, April 2001. `http://www.ripe.net/ttm`.

[13] R. Hinden. New scheme for internet routing and addressing (ENCAPS) for IPNG. *RFC 1955*, June 1996.

[14] G. Huston. The growth of the bgp table - 1994 to present. Available online at: `http://bgp.potaroo.net`.

[15] D. Massey, L. Wang, B. Zhang, and L. Zhang. A proposal for scalable internet routing & addressing. *Internet Draft*, February 2007. Available online at: `www.ietf.org/internet-drafts/draft-wang-ietf-efit-00.txt`.

[16] D. Meyer. University of Oregon Route Views Project. `http://www.routeviews.org`, January 2005.

[17] D. Meyer, L. Zhang, and K. Fall. Report from the IAB workshop on routing and adressing. Available online at: `www.ietf.org/internet-drafts/draft-iab-raws-report-01.txt`.

[18] R. Moskowitz and P. Nikander. Host identity protocol (HIP) architecture. *RFC 4423*, May 2006.

[19] E. Nordmark and M. Bagnulo. Level 3 multihoming shim protocol. *Internet Draft*, May 2006. Available online at: `www.ietf.org/internet-drafts/draft-ietf-shim6-proto-07.txt`.

[20] M. O'Dell. GSE - an alternate addressing architecture for IPv6. *Internet Draft*, 1997. `http://www.watersprings.org/pub/id/draft-ietf-ipngwg-gseaddr-00.txt`.

[21] B. Quoitin. *BGP-based Interdomain Traffic Engineering*. Université Catholique de Louvain, August 2006. PhD Thesis.

[22] B. Quoitin and O. Bonaventure. A Cooperative Approach to Interdomain Traffic Engineering. In *Proceedings of EuroNGI*, April 2005.

[23] RIPE. Routing Information Services. `http://www.ripe.net/projects/ris/index.html`, March 2007.

[24] J. Saltzer. On the naming and binding of network destinations. *RFC 1498*, August 1993.

[25] I. Stoica, D. Adkins, S. Zhuang, S. Shenker, and S. Surana. Internet Indirection Infrastructure. In *Proceedings of ACM SIGCOMM*, August 2002.

[26] X. Yang. NIRA: a new internet routing architecture. *Proceedings of the ACM SIGCOMM workshop on Future directions in network architecture FDNA '03*, August 2003.

[27] H. Zhang and M. Chen. Forming an IPv6-only Core for Today's Internet. In *Proceedings of ACM SIGCOMM IPv6 Workshop*, August 2007.

[28] X. Zhang, P. Francis, J. Wang, and K. Yoshida. Scaling global IP routing with the core router-integrated overlay. *Proceedings of the 14th IEEE International Conference on Network Protocols*, November 2006.