Locator/Identifier Split

Wiebke Köpp
Advisor: Alexander Klein
Seminar Future Internet SS2011
Chair for Network Architectures and Services
Department of Computer Science, Technical University of Munich
Email: koepp@in.tum.de

ABSTRACT

The size of routing tables in the default free zone (DFZ) has exceeded 350000 entries by now and will grow even more in the future. The reasons behind this rapid growth are provider-independent addressing, multihoming and traffic engineering. IPv6, providing a much bigger address space than IPv4, allows for more devices to be connected. As a consequence, routing does not scale anymore and measures have to be taken in order to reconstitute scalability in the Internet. Many approaches which try to do that are based on a Locator/Identifier (Loc/ID) split. It modifies the current addressing paradigm by splitting locators for routing purposes from identifiers of end-systems. This paper is a survey of the Loc/ID split. It explains its general ideas and describes implementation efforts.

Keywords

Locator/Identifier Split, Routing, Scalability, LISP

1. INTRODUCTION

In its early stages, the Internet was only a network of a few research facilities. Nowadays over 2 billion people have access to the Internet, with even more to come [1]. The number of hosts has reached a point where they cannot be numbered using the address space of IPv4. IPv6 provides a much bigger address space, making it possible to connect more users and devices, as a result solving the issue of address depletion. However, IPv6 reinforces scalability problems at the same time. Besides the increase of users, different behaviors of Internet Service Providers (ISPs) and their customers challenge the current Internet architecture. There has been a shift in how customers use the Internet. A growing interest in multihoming, thus being connected to multiple providers instead of just one, can be recognized since users want reliable access to the Internet at all time. Also, an increasing number of mobile devices are connected to the Internet, creating a demand for support of mobility. Providers on the other hand perform traffic engineering. The way these actions are performed and these demands are fulfilled today are reasons for the rapid growth of routing tables in the DFZ, as shown in Figure 1. The present routing architecture will not be able to scale having to cope with the resulting entries. Another reason for the observed scalability problems are the overloaded semantics of IP addresses [27]. IP addresses are used for both identifying end-systems and locating them for routing purposes. Yakov Rekhter once stated: *"Addressing follows topology or topology follows Addressing. Choose one."* [27] But routing is most efficient when addresses are assigned topologically, while handling of end-systems requires exactly the opposite. The single numbering space currently in use certainly cannot serve both. Therefore, a split into two separate spaces, one for identifiers and one for locators has been proposed. Using Loc/Id split principles a host has an identifier and a locator instead of one address for both purposes. Some of the Loc/ID split proposals mainly focus on mobility, but other approaches are expected to solve all the issues addressed above.

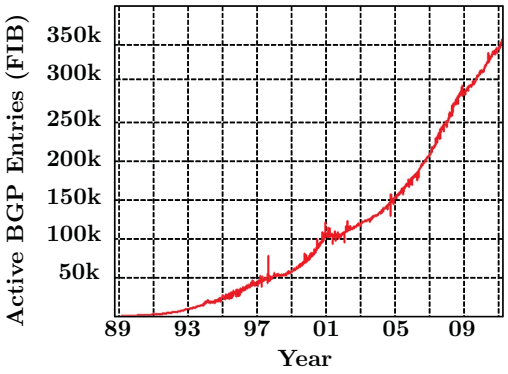

Figure 1: Size of routing tables in the DFZ [14]

This paper is structured as follows. A motivation for the Loc/ID split and a short review of current routing is given in Section 2. General ideas and performance measurements are explained in Section 3. Section 4 introduces one specific implementation in detail and gives a brief overview of other approaches. Section 5 finally concludes the paper.

2. MOTIVATION FOR THE SPLIT

As already mentioned in the introduction, today's routing does not scale anymore for various reasons. Accordingly, some changes to the inter-domain routing protocol Border Gateway Protocol (BGP) have been proposed (e.g. [2, 31, 7]). Unfortunately they imply changes that are hard to deploy.

2.1 State-of-the-art routing and its problems

2.1.1 Current routing

The Internet consists of over 35000 autonomous systems (ASes) [14]. Communication can happen either within a single AS or between multiple ASes. Therefore, a distinction

doi: 10.2313/NET-2011-07-2_09

between intra- and inter-domain routing is drawn. Intra-domain routing takes place within a single AS, while Inter-domain routing handles communication between different ASes. Both use different routing protocols that are adapted to the location where they are deployed.

Routing tables within an AS are created by assigning administrative costs to all links and then use the path with the lowest cost to forward traffic. Intra-domain routing mostly uses link-state routing protocols like OSPF [30] or IS-IS [29]. In link-state routing protocols, routers do not only learn who their neighbors are, they also receive the topology of the system. Using this information, routers can calculate the best path to a desired destination. In intra-domain routing an additional default route can be specified if routing tables contain no match to the destination address. Larger ASes sometimes divide their network into several smaller networks to keep routing within their AS simple and scalable.

Inter-domain routing, on the other hand, uses BGP, which is a path vector protocol. A BGP router tells its neighbors which prefixes are reachable over its own network and which ASes need to be traversed to reach a destination AS. A router then looks up a packet's next hop by searching for the longest prefix match in its forwarding information base (FIB) which is based on the information found in routing tables. Contrary to routers in edge networks, routers in the DFZ do not provide a default route if no match for the destination address can be found. They have an entry for each reachable prefix, causing routing tables to grow with every additional reachable prefix [23].

2.1.2 Scalability

The number of entries in routing tables in the DFZ is increasing rapidly. Along with routing table size, update rates also rise. Update rates are usually around 1–10 update messages per second and peak at approximately 1000 updates per second. Additionally, with the transition to IPv6, larger update rates and routing tables can be expected. As a result, future routers must be very powerful in order to answer all route requests without significant delay. They need to process traffic fast, handle a large amount of updates and store all needed information in their memory. Researchers argue this cannot be accomplished at reasonable cost [27].

Thus, handling the growth of routing tables can only be achieved by eliminating the reasons for the growth. The reasons are provider-independent addressing, multihoming, traffic engineering and countermeasures against prefix hijacking.

In general, IP address space can be owned by either providers or customers. If the provider is the owner of the address space and the customer only rents it, the addresses are called provider-aggregatable (PA). If the addresses belong to the customer, they are provider-independent (PI). A customer with PA addresses has to renumber all his devices when changing providers. Since his new IP address space is a subspace of the providers AS, no additional entries or BGP updates are needed. However, the renumbering is still a costly process which customers would rather avoid. In consequence many customers prefer PI addresses. Provider changes of customers with PI addresses cause updates and new BGP entries because their addresses are usually not aggregatable with the ones of the new provider.

Customers can be interested in multihoming for different purposes. Reliability and service differentiation are two examples. The connection to multiple ISPs can make the Internet connection of a customer more reliable. If the connection to one ISP fails, a fallback connection to another provider can be used. This causes several BGP entries for a single prefix. A customer could also decide to use different providers for diverse services. A portion of the customer's network is assigned to each service, which is then connected to an ISP. Several longer prefixes are announced to BGP instead of one prefix for the whole network.

Providers use traffic engineering to improve performance and use their network's resources more efficiently. For example, they announce more specific routes, thus longer prefixes, to BGP in order to attract traffic at certain gateways. Countermeasures against prefix hijacking also cause providers to announce long prefixes into BGP. In this case, the longest possible prefix is injected into BGP, to prevent a malicious AS to insert a longer prefix and thereby attracting all the traffic. This is mostly done for important services like the Domain Name System (DNS) [27, 23].

2.1.3 Mobility

Researchers expect the number of mobile Internet users to surpass the number of fixed Internet users by 2014 [12]. They also state that with use of the Loc/ID split the impact on routing scalability could be kept at a minimum. The main challenge in mobility is to maintain the connection between hosts, e.g. TCP/IP connections, even if one of the hosts is changing its location. TCP uses IP addresses as identifiers for a connection. When a host moves from one network to another, his IP address changes and the connection is lost. For this reason, the IP address should stay the same, but this interferes with Internet routing. The Loc/ID split provides a solution to accomplish both, keeping IP address and still be able to route properly at the same time.

2.2 Proposed Enhancements to BGP

Efforts have been made to make BGP scalable again, but those are almost impossible to deploy since the Internet is too widely distributed to swap out a protocol on a certain day. Other proposals leave BGP as it is today and introduce an overlaying architecture. Among these are aggregation proxies [34] and lookup systems for nonroutable prefixes [13]. Unfortunately, both approaches are hard to deploy since they require major changes to the Internet in order to work efficiently [23].

2.2.1 Aggregation proxy

ISPs do not announce their prefixes directly to BGP, but to an aggregation proxy. The proxy receives multiple long prefixes, aggregates them to a shorter prefix and then announces the result to BGP. Traffic directed to one of these ISPs is always routed via the proxy, which tunnels the packets to the right destination. The example in Figure 2 shows four networks with prefixes `10.10.0.0/24`, `10.10.1.0/24`, `10.10.2.0/24` and `10.10.3.0/24`. The aggregation proxy aggregates the networks to the shorter prefix `10.10.0.0/22`

doi: 10.2313/NET-2011-07-2_09

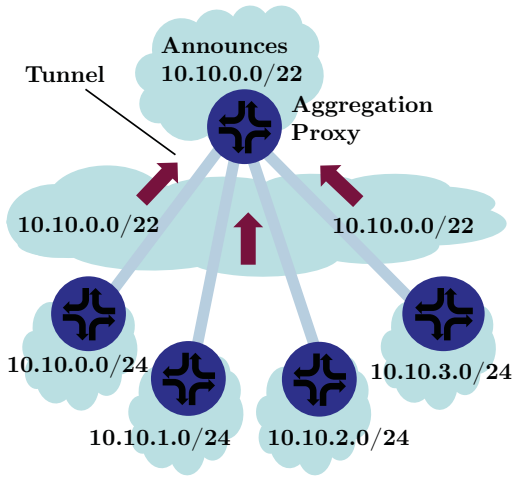

Figure 2: Aggregation Proxy, see [23]

and inserts it into BGP routing tables. This way, routing tables altogether contain fewer entries. In the example, the routing table size is reduced by 3. However, routing via the proxy could lead to longer paths compared to the original BGP path. Another disadvantage is that it is not clear who should be in charge of operating proxies [23].

2.2.2 Lookup System for nonroutable prefixes

Similar to the concept of aggregation proxies, long prefixes are not announced, when using a lookup system for nonroutable prefixes. Instead, they are put in a DNS-like lookup system. Along with the prefix, an entry contains a router over which the prefix is reachable. This router is usually part of the same AS as the prefix. The prefixes in the lookup system do not occur in BGP routing tables, thus they are not routable in the DFZ. If a router receives a packet he cannot find a matching prefix for, he queries the lookup system. The lookup system replies with the address of the router the destination address can be reached over. The router then encapsulates the packet towards the received address, where it is decapsulated again and forwarded to the destination address via intra-domain routing. The process is shown in Figure 3.

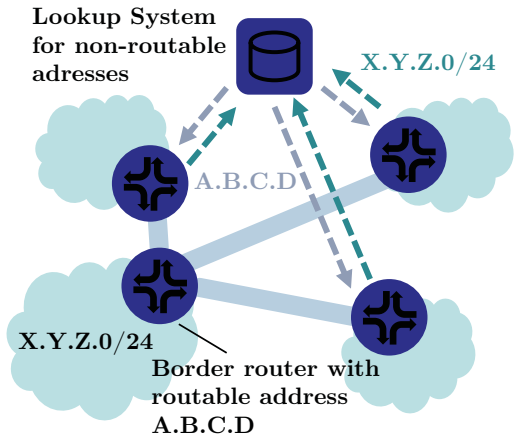

Figure 3: Lookup System, see [23]

Deploying this approach requires even more reformation than the deployment of aggregation proxies. The functionality of looking up nonroutable prefixes on one hand and tunneling packets on the other hand has to be introduced in BGP routers. Additionally, the lookup system itself needs to be created [23].

3. THE LOCATOR/IDENTIFIER SPLIT

The Loc/ID split is a principle many proposals solving scalability issues use. The following section first describes general ideas different approaches have in common and then analyses the performance of those approaches.

3.1 General Ideas

All Loc/ID split solutions have in common that they create two different namespaces for locators and identifiers. While in some approaches both locator and identifier remain IPv4 or IPv6 addresses, other solutions create a new namespace for identifiers.

The Loc/ID solutions that have been proposed so far, fall into two major categories: Map-and-encap and Address Rewriting. [6] gives a detailed comparison of both approaches that will be briefly explained in the following. They can also be classified by the network element, at which they require changes. Host-based solutions require changes at hosts while router-based solutions imply new functionalities in routers. Hybrid solutions also exist.

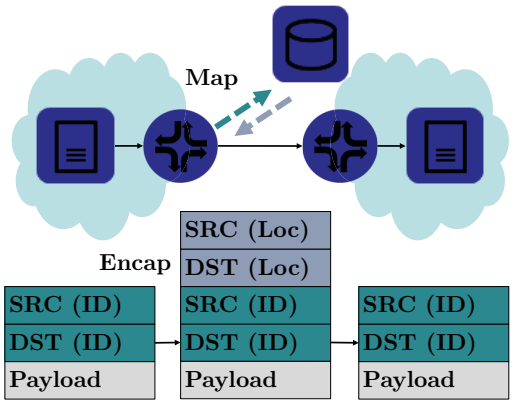

Figure 4: Map and Encap

Map-and-encap stands for approaches that use a mapping system and encapsulation. A host that wants to send data to another host outside its own domain starts by looking up the IP-address of the desired destination in DNS and fills it into the IP-header. Then, the packet travels through the AS. The border router looks up the locator for the destination address and encapsulates the packet. Next, BGP is used to transport the data to the router with that address, where it is decapsulated again and then forwarded to the destination host. The principle can be seen in Figure 4. Map-and-encap solves discussed challenges for scalability issues: Customers can easily switch providers because instead of BGP updates and new entries, only the locator-identifier-mapping in the mapping systems has to be updated. Multihoming and traffic engineering are also supported. The mapping-system can contain several locators for an identi-

doi: 10.2313/NET-2011-07-2_09

fier. Inside an entry in the mapping system, priorities and weights can be assigned to a locator, so that traffic can be directed in a desired way. However, encapsulation adds an additional header to a packet. Packet sizes might come in conflict with Maximum Transfer Units (MTU) and require fragmentation in consequence [23, 25].

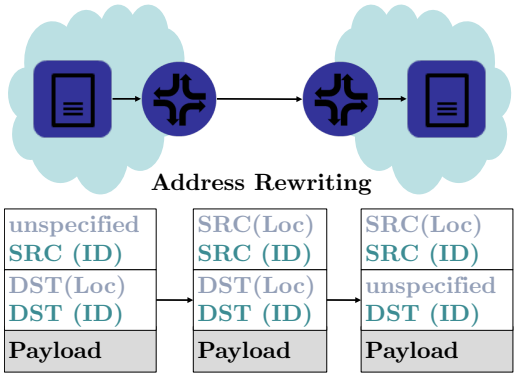

Figure 5: Address Rewriting

Address rewriting solutions take advantage of 128-bit IPv6 address, using the top 64 bits as a locator and the lower 64 bits as an identifier. A host sending a packet specifies the lower 64 bits by inserting its identifier. What happens to the top 64 bits differs in the individual implementations. The top 64 bits could for example be filled with an unspecified value, when a host has no information about its corresponding locator. The locator is then filled in by the border router from where the packet is traversed to the border router of the destination AS. This router replaces the destination locator bits and directs the packet to the destination host [25]. The basic concept can be seen in Figure 5. While some address rewriting approaches require a mapping system, other assume certain host abilities. Once a border router has made an initial address choice, the host is supposed to use that choice in ongoing communications. Multihoming and traffic engineering are also controlled by border routers. They can change the source address of outgoing data in order to redirect returning traffic. When changing providers there is again no need for renumbering since the global routing architecture has no knowledge of identifiers. [6]

In the area of mapping systems used by map-and-encap approaches, several things need to be considered. One goal is to keep the product $State \times Rate$ small. Rate refers to the update rate of identifier-locator-mappings. State means the size of the mapping system in bits. Since most estimates put state around $\mathcal{O}(10^{10})$ [26], the update rate should be small. The same accounts for latency triggered by the lookup in the mapping system. Since a router has to query the mapping system every time it has no information about an identifier, it is convenient for routers to have a local cache. In some mapping system proposals, a router even holds a copy of the whole mapping database. If a mapping is not in the cache, a packet can either be stored and delayed, dropped or forwarded to a place where the mapping is known. Obviously this should occur very infrequently. There exist three different kinds of mapping systems: Pull, Push and hybrid systems that apply a push/pull strategy. In pull systems the router is responsible to maintain mapping entries, while in a push model the mapping service itself initiates updates. Hybrid systems push only some of the data, for example to intermediate databases, while others mappings need to be specifically pulled [23, 25].

Another issue to be considered when implementing a new protocol is incremental deployability. That means protocols should always provide ways to interoperate with the legacy Internet [18]. Otherwise, a protocol has to be deployed in a widely manner from a specific day on to benefit the Internet's architecture. Some approaches achieve incremental deployability by introducing additional proxy gateways, others do not need additional entities.

3.2 Performance Analysis

In terms of scalability, the Loc/ID split has two major implications on inter-domain routing. First, BGP routing tables are reduced and second, the BGP update rate decreases. Using data from a time span from January 2004 until June 2008, Dong et al. [8] could make assumptions on how big the impact of Loc/ID split would be. They came to the following results. There are two different kinds of ASes: stub ASes and transit ASes. Transit ASes deliver data to other ASes and correspond to service providers. Stub ASes, on the other hand, only appear at the end of an ASes path, hence they are customer networks. Stub ASes account for about 80 percent of the total number of ASes, leaving around 20 percent for transit ASes. Even though transit ASes take up only about 20 percent of the number of ASes, the fraction of prefixes belonging to them is much bigger with 60 to 65 percent. This is because transit ASes usually are larger than stub ASes. The Loc/ID split keeps multihoming and traffic engineering activity in customer networks away from routing in the DFZ. Prefixes of customer networks are not announced into BGP. Thus, routing tables can be reduced by the number of prefixes in stub ASes. Dong et. al. also state that stub ASes are responsible for approximately 50 to 60 percent off the updates in BGP which would be eliminated with Loc/ID deployment. BGP update rates are thereby also reduced.

Map-and-encap protocols add an additional header. On account of this header, traffic overhead is produced. Iannone and Bonaventure measured this overhead among other things using data collected at their university and published their results in [15]. They state that 4 to 15 percent overhead in outgoing and 2 to 10 percent in incoming traffic are produced. In their opinion overhead caused by encapsulation should therefore not cause any problems. In [16], the authors report that the additional delay due to encapsulation and decapsulation packet forwarding is around $1\mu s$, only decapsulation in IPv6 causes higher delay with approximately $3\mu s$. The extent to which encapsulation causes problems with MTU is to date unclear, but some possibilities to deal with MTU issues are suggested in [9].

The performance of many Loc/ID solutions additionally depends on the mapping system. In general, in push systems latency is small, while state is big. With pull systems it is the other way around. In push systems, routers already have all the mapping information on cost of having to save it. Pull systems require less memory, but have to query the mapping system more often instead. Hybrid systems create

doi: 10.2313/NET-2011-07-2_09

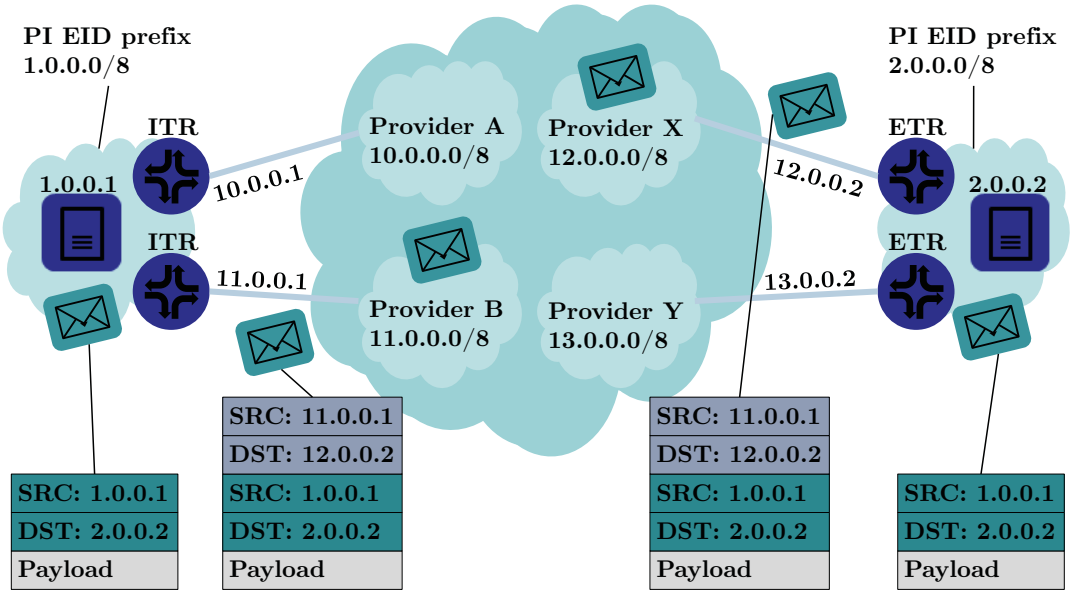

Figure 6: Path of a packet in LISP, see [26]

a balance between state and latency. The state of databases in mapping systems has been addressed in [22]. Assuming that the average identifier multihomes with three ISPs and additional information for, e.g. security, is saved in an entry, the authors of this paper claim that the entry in a mapping database when using IPv6 is about 178 bytes long. With currently about 10^6 needed entries, mapping databases take up about 178 MB. Evaluation of performance in mapping systems in terms of caching hit rates and update rates have also been discussed in several papers [22, 15, 33]. But they each focus on one specific approach and are, therefore, not discussed here.

4. IMPLEMENTATION

This section covers specific Loc/ID split solutions. First, the Locator Identifier Separation Protocol (LISP) is explained. Since LISP-Alternative-Topology (LISP+ALT) is the currently preferred mapping system with LISP, it is discussed as well. Finally, a brief overview on other Loc/ID protocols is given.

4.1 Locator Identifier Separation Protocol

LISP is a map-and-encap-protocol. The split is done by introducing endpoint identifiers (EIDs) and routing locators (RLOCs). Both are IPv4 or IPv6 addresses, but only RLOCs are routable in BGP. The border router performing encapsulation and decapsulation are called ingress tunnel routers (ITR) and egress tunnel routers (ETR) [9]. The communication between two hosts in the same LISP-domain works exactly as it does today. But, if two hosts in different LISP-domains want to communicate, the map-and-encap mechanism is needed. For this, the mapping system is required. It is queried by an ITR, when the ITR does not have a so-called EID-to-RLOC mapping in its local cache. An ITR is responsible for encapsulating a package and sending it towards the specified ETR in the destination domain. An ETR on the other hand has to decapsulate incoming traffic and forward data to its destination.

An example for the path of a packet can be seen in Figure 6. The steps from the first host sending the packet until the second host receives it are the following [26]:

1. The host with EID `1.0.0.1` wants to send a packet to the host with EID `2.0.0.2`. It simply puts those addresses in an IP packet and sends it.
2. The packet traverses the AS until it reaches a border router. In this case, it is the ITR with RLOC `11.0.0.1`
3. Assuming the ITR already has a mapping for `2.0.0.2` in its cache, it encapsulates the packet with a new header. Here, the mapping of EID `2.0.0.2` returns the RLOC `12.0.0.2`. `13.0.0.2` also could have been chosen. The additional header has the RLOC of the ITR as source address and the result of the EID-to-RLOC-mapping of `2.0.0.2` as destination address.
4. Next, the data is send to the ETR with RLOC `12.0.0.2` using BGP.
5. The ETR decapsulates the packet and forwards the packet to the destination address.

As already mentioned, protocols should provide some way to be reachable by hosts in the legacy Internet and also reach those nodes themselves. When a host wants to send packets to a non-LISP host, the ITR could simply forward the packet without encapsulation. But most providers make sure they do not process traffic not belonging to their customers. As a solution, a proxy ETR (PETR) and a proxy ITR (PITR) have been introduced. They are located in networks that do not check source addresses that way. The IP address of a legacy host is not listed in the mapping system with EID-to-RLOC-mappings. So, if an ITR cannot find a mapping for

doi: 10.2313/NET-2011-07-2_09

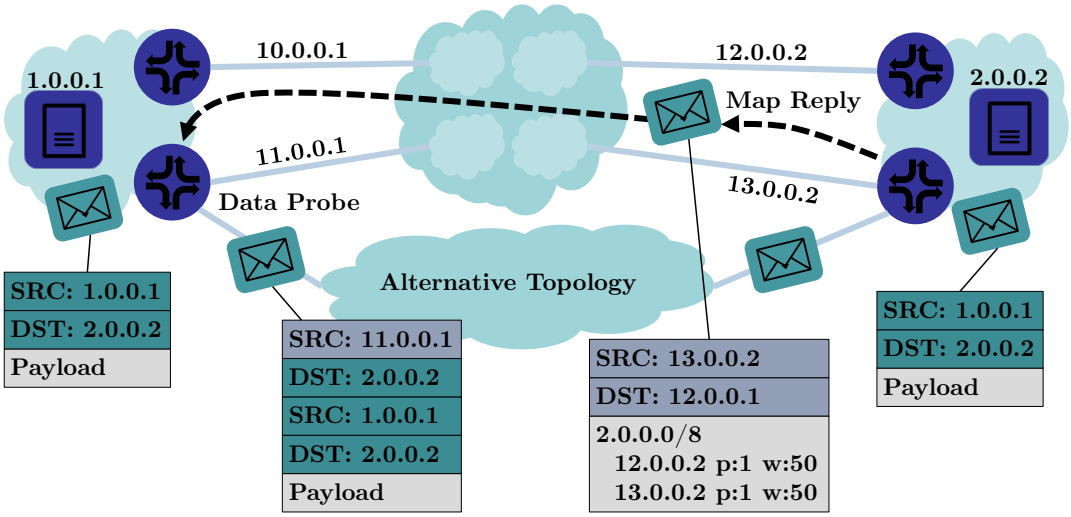

Figure 7: Path of a packet in LISP+ALT when a cache miss occurs, see [26]

an EID it detects that communication with a legacy host is wished. Then, the ITR forwards the data to a corresponding PETR, which can simple transport it to the destination. In case of the other direction, PITRs announce all prefixes via anycast, they want to attract traffic for, thus all addresses used for EIDs. This has to be done since EIDs are not globally routable. PITRs then basically work like normal ITRs. They perform a mapping and then tunnel data to their destination [24].

LISP is a protocol that also provides an architecture to integrate mobile hosts. Its name is LISP Mobile Node (LISP MN). LISP MN allows mobile hosts to multihome and TCP connections to stay alive while roaming. The current proposal of LISP MN seems to have some advantages over Mobile IP, the most common approach to provide mobility, but new problems have also been experienced. A detailed description is given in [24]. The problems are also addressed in this article.

4.2 LISP Alternative Topology

LISP Alternative Topology (LISP+ALT) [10] is the currently preferred mapping systems used with LISP. Basically this mapping system introduces an overlay structure only for handling EID-to-RLOC-mappings. LISP+ALT is a hybrid push/pull model. In case the whole database is pushed to LISP+ALT routers and ITRs have a cache pulling mappings as needed. In case of an ITR needing to get a specific mapping, thus when a cache miss occurs, the ITR can send either a Mapping Request or a Data Probe. Both are routed over the alternative topology and result in a Map Reply sent by the ETR of the original destination EID. A Data Probe is a special kind of Map Request which already contains the data to be sent. This prevents dropping or delaying packets at the ITR. The alternative topology then uses BGP to forward the data. In order to send a data probe an ITR has to encapsulate the packet. Since the destination RLOC is unknown, the destination EID is simply copied.

Figure 7 shows a scenario when a cache miss is experienced. Here, a data probe is used to retrieve the mapping. The following steps are performed.

1. The host with EID `1.0.0.1` wants to send a packet to the host with EID `2.0.0.2`.
2. The packet traverses the AS until it reaches a border router. In this case, it is the ITR with RLOC `11.0.01`
3. This time the ITR does not have a mapping for `2.0.0.2` in its cache. It therefore encapsulates the packet with a new header and sends it to the mapping system. Since the ITR has no knowledge of the destination RLOC it fills the destination EID in the header instead.
4. The mapping system then takes care of transporting the package to the destination RLOC. In this example `13.0.0.2` is used.
5. The ETR decapsulates the packet and forwards the packet to the destination address. It also sends a map reply to back to `11.0.0.1` telling this ITR that the prefix `2.0.0.0/8` can be reached over the two RLOCs `12.0.0.2` and `13.0.0.2`. Additionally a priority (p) and a weight (w) for each RLOC are provided. Here, both RLOCS have the same priority and weight, so that traffic should be split equally among both.

Researchers proposed many other approaches for mapping systems, namely: APT [20], CONS [4], DHT [19], EMACS [5], FRMS [22] and NERD [17]. Table 1 shows which distribution model is used on each case. A brief description of these mapping systems can be found in [21].

4.3 Other Approaches

The Host Identity Protocol (HIP) is a Loc/ID split protocol that adds an additional layer between transport and network layer in the OSI model. IP addresses are kept as locators, but identifiers get a completely new namespace. The so-called Host Identifier (HI) is a cryptographic public key. HIs are usually not used directly. Instead, a Host

doi: 10.2313/NET-2011-07-2_09

Table 1: Distribution Models [19]

Mapping system	Distributions Model
ALT	Hybrid Push/Pull
APT	Push
CONS	Hybrid Push/Pull
DHT	Pull
EMACS	Pull
FRMS	Hybrid Push/Pull
NERD	Push

Identifier Tag (HIT) is used to represent the HI. The HIT is a hash of the HI and has a fixed length of 128 bit, exactly the size of IPv6 addresses. A representation having the size of an IPv4 address also exists, the Local Scope Identifier (LSI). These sizes are used for support of legacy applications. HITs or LSIs replace IP addresses as identifiers in TCP connections. In order to establish a connection a four-way handshake between two hosts, called Initiator and Responder, is performed. HIP is expected to improve mobility and make multihoming easier. Furthermore, connections over HIP are more secure due to the use of public keys as identifiers. Therefore, HIP combines several functionalities that are usually provided by separate protocols [28].

The Shim6 architecture is a Loc/Id approach dealing mostly with multihoming. Both Shim6 and HIP are host-based. The Shim6 architecture introduces a new layer like HIP and two new protocols: Shim6 and the reachability protocol (REAP). Shim6 is the protocol establishing a connection between two hosts, thus creating a Shim6 context. Shim6 also uses a four-way handshake. During that handshake a set of locators for the two identifiers, called upper-layer identifiers (ULID) in Shim6, is exchanged. A different context can be used for each direction. During ongoing connections a host can send an Update Request containing a new set of locators, which is then answered by an Update Acknowledgment. REAP is a protocol in charge of failure detection. A communication usually consists of data traffic in both directions. If there is only traffic in one direction, REAP will send keepalives in the other direction. If at some point there is no incoming traffic at either one of the hosts, a failure is assumed [11].

Six/One Router is an address rewriting approach acting at the router. A network deploying Six/One Router usually consists of so-called PI edge addresses. Additionally the network is assigned a set of transit addresses by each of its providers. One edge address can be mapped on exactly one transit address per provider and one transit address corresponds to one edge address. Six/One Router uses a mapping system. Either one of the mapping systems proposed for LISP can be used for that purpose. Border routers, also called Six/One Routers are responsible for translating edge to transit addresses. Edge addresses in Six/One Router networks are not routable in the DFZ directly, they can only be reached through their transit addresses. Each time a packet crosses the border of an edge network a mapping has to be performed. For incoming packets, the destination address has to be translated into an edge address. For outgoing data the source address has to be modified. Six/One Router is a protocol designed to solve multihoming issues [32].

Another address rewriting approach is the Identifier-Locator Network Protocol (ILNP). It is also a host-based solution. ILNP uses the same packet format as IPv6, but splits source and destination address in half. 64 bits are used as a locator and 64 bits are used as an identifier. In this protocol, the identifiers are encoded MAC-addresses. Locators specify a subnet. ILNP works with DNS, where new kinds of entries need to be created [3].

5. CONCLUSION

The Loc/ID split is a principle expected to overcome scalability issues in the current Internet routing architecture while maintaining efficient support for multihoming, traffic engineering and PI addresses. The performance analysis confirmed that routing tables in the DFZ as well as the BGP update rate can be significantly reduced by deploying a Loc/ID principle. 35 to 40 percent of the prefixes in routing tables can be eliminated by deploying a Loc/Id principle. Updates can be reduced by 50 to 60 percent. Many different proposals for new protocols using a Loc/ID approach exist. This paper described LISP as an example of an implementation mainly focusing on scalability. In the following a few other approaches have been mentioned. These are HIP, Shim6, Six/One Router and ILNP. HIP, Shim6 and ILNP are host-based solutions, which means that they require changes to the host. Six/One Router and LISP are router-based protocols and basically require no changes to hosts. While LISP is a protocol mostly focusing on scalability issues, HIP for example addresses secure mobility. Thus, the Loc/Id split can contribute to solving different issues. LISP and some of the other approaches require an additional mapping system to map identifiers to locators. LISP+ALT has been discussed in order to give such an example. In general, mapping systems should minimize the product of update rate and size of the mapping system ($State \times Rate$). Since the number of hosts already is high and will grow further in the future, the update rate in the mapping system should be kept low. Performance in mapping systems has been discussed in several papers, but still requires further investigation. This should be done to show that mapping systems are really efficient and are not only shifting scalability problems from the routing architecture to mapping systems. The new scalable routing architecture relies on a mapping system. If that mapping system does not scale no progress will be made by introducing it.

6. REFERENCES

[1] ITU Statistics. `http://www.itu.int/ITU-D/ict/statistics/`, Mar. 2011 (accessed March 28, 2011).

[2] Y. Afek, A. Bremler-Barr, and S. Schwarz. Improved BGP Convergence via Ghost Flushing. *IEEE Journal on Selected Areas in Communications*, 22(10):1933–1948, 2004.

[3] R. Atkinson. An Overview of the Identifier-Locator Network Protocol (ILNP). Research Note RN/05/26, University College London, Sept. 2005.

[4] S. Brim, N. Chiappa, D. Farinacci, V. Fuller, D. Lewis, and D. Meyer. LISP-CONS: A Content distribution Overlay Network Service for LISP `http://tools.ietf.org/html/draft-meyer-lisp-cons-04`, IETF, Apr. 2008. Work in

doi: 10.2313/NET-2011-07-2_09

progress.

[5] S. Brim, D. Farinacci, D. Meyer, and J. Curran. EID Mappings Multicast Across Cooperating Systems for LISP. http://tools.ietf.org/html/draft-curran-lisp-emacs-00, IETF, Nov. 2007. Work in progress.

[6] L. Burness, P. Eardley, S. Jiang, and X. Xu. A pragmatic comparison of locator ID split solutions for routing system scalability. In *Third International Conference on Communications and Networking in China, 2008. ChinaCom 2008.*, pages 1024–1028, 2008.

[7] J. Chandrashekar, Z. Duan, Z.-L. Zhang, and J. Krasky. Limiting Path Exploration in BGP. In *INFOCOM 2005*, volume 4, pages 2337–2348, 2005.

[8] P. Dong, H. Wang, Y. Qin, H. Zhang, and S.-Y. Kuo. Evaluation of Scalable Routing Architecture Based on Locator/Identifier Separation. In *GLOBECOM Workshops, 2009 IEEE*, pages 1–6, Dec. 2009.

[9] D. Farinacci, V. Fuller, D. Meyer, and D. Lewis. Locator/ID Separation Protocol (LISP). http://www.ietf.org/id/draft-ietf-lisp-11.txt, IETF, Mar. 2011. Work in progress.

[10] V. Fuller, D. Farinacci, D. Meyer, and D. Lewis. LISP Alternative Topology (LISP+ALT). http://tools.ietf.org/html/draft-ietf-lisp-alt-06, IETF, Mar. 2011. Work in progress.

[11] A. Garcia-Martinez, M. Bagnulo, and I. Van Beijnum. The Shim6 Architecture for IPv6 Multihoming. *IEEE Communications Magazine*, 48(9):152–157, 2010.

[12] M. Grayson, K. Shatzkamer, and K. Wierenga. *Building the Mobile Internet*. Cisco Press, Feb. 2011.

[13] W. Herrin. Tunneling Route Reduction Protocol (TRRP). http://bill.herrin.us/network/trrp.html, 2008 (accessed March 31, 2011).

[14] G. Houston. The BGP Routing Table. http://bgp.potaroo.net/, Mar. 2011 (accessed March 25, 2011).

[15] L. Iannone and O. Bonaventure. On the Cost of Caching Locator/ID Mappings. In *Proceedings of the 2007 ACM CoNEXT conference*, pages 7:1–7:12. ACM, 2007.

[16] L. Iannone, D. Saucez, and O. Bonaventure. Implementing the Locator/ID Separation Protocol: Design and experience. *Computer Networks*, 55(4):948–958, Mar. 2011.

[17] E. Lear. NERD: A Not-so-novel EID to RLOC Database. http://tools.ietf.org/html/draft-lear-lisp-nerd-08, IETF, Mar. 2010. Work in progress.

[18] T. Li. Design Goals for Scalable Internet Routing. http://www.ietf.org/id/draft-irtf-rrg-design-goals-06.txt, IETF, Jan. 2011. Work in progress.

[19] L. Mathy and L. Iannone. LISP-DHT: Towards a DHT to map Identifiers onto Locators. In *Proceedings of the 2008 ACM CoNEXT Conference*, CoNEXT '08, pages 61:1–61:6, 2008.

[20] M. Meisel, D. Massey, L. Wang, B. Zhang, and L. Zhang. APT: A Practical Transit Mapping Service. http://tools.ietf.org/html/draft-jen-apt-01, IETF, Nov. 2007. Work in progress.

[21] M. Menth, M. Hartmann, and M. Hoefling. Mapping Systems for Loc/ID Split Internet Routing. Technical Report 472, University of Würzburg, May 2010.

[22] M. Menth, M. Hartmann, and M. Hofling. FIRMS: A Mapping System for Future Internet Routing. *IEEE Journal on Selected Areas in Communications*, 28(8):1326–1331, 2010.

[23] M. Menth, M. Hartmann, P. Tran-Gia, and D. Klein. Future Internet Routing: Motivation and Design Issues. *it - Information Technology*, 50(6):358–375, 2008.

[24] M. Menth, D. Klein, and M. Hartmann. Improvements to LISP Mobile Node. In *ITC 22nd International Teletraffic Congress (ITC22), Amsterdam 2010*, pages 1–8, 2010.

[25] D. Meyer. Update on Routing and Addressing at IETF 69. *IETF Journal*, 3(2):21–24, Oct. 2007.

[26] D. Meyer. The Locator Identifier Separation Protocol (LISP). *The Internet Protocol Journal*, 11(1):23–36, 2008.

[27] D. Meyer, L. Zhang, and K. Fall. Report from the IAB Workshop on Routing and Addressing. RFC 4984, http://www.ietf.org/rfc/rfc4984.txt, Sept. 2007.

[28] P. Nikander, A. Gurtov, and T. Henderson. Host Identity Protocol (HIP): Connectivity, Mobility, Multi-Homing, Security, and Privacy over IPv4 and IPv6 networks. *Communications Surveys Tutorials, IEEE*, 12(2):186–204, 2010.

[29] D. Oran. OSI IS-IS Intra-domain Routing Protocol. RFC 1142, http://www.ietf.org/rfc/rfc1142.txt, Feb. 1990.

[30] J. Oran. OSPF Version 2. RFC 2328, http://www.ietf.org/rfc/rfc2328.txt, Apr. 1998.

[31] W. Sun, Z. Mao, and K. Shin. Differentiated BGP update processing for improved routing convergence. In *Proceedings of the 14th IEEE International Conference on Network Protocols, 2006. ICNP '06.*, pages 280–289, 2006.

[32] C. Vogt. Six/One Router: A Scalable and Backwards Compatible Solution for Provider-Independent Addressing. In *Proceedings of the 3rd international workshop on Mobility in the evolving internet architecture*, MobiArch '08, pages 13–18, 2008.

[33] H. Zhang, M. Chen, and Y. Zhu. Evaluating the performance on Id/Loc mapping. In *IEEE GLOBECOM 2008*, pages 1–5, 2008.

[34] X. Zhang, P. Francis, J. Wang, and K. Yoshida. Scaling IP routing with the Core Router-integrated overlay. In *Proceedings of the 2006 14th IEEE International Conference on Network Protocols, 2006. ICNP '06.*, pages 147–156, 2006.

doi: 10.2313/NET-2011-07-2_09