Internet Engineering Task Force (IETF)
Request for Comments: 7834
Category: Informational
ISSN: 2070-1721

D. Saucez
INRIA
L. Iannone
Telecom ParisTech
A. Cabellos
F. Coras
Technical University of Catalonia

April 2016

Locator/ID Separation Protocol (LISP) Impact

Abstract

The Locator/ID Separation Protocol (LISP) aims to improve the Internet routing scalability properties by leveraging three principles: address role separation, encapsulation, and mapping. In this document, based on implementation work, deployment experiences, and theoretical studies, we discuss the impact that the deployment of LISP can have on both the routing infrastructure and the end user.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7834.

Table of Contents

1. Introduction

The Locator/ID Separation Protocol (LISP) relies on three principles to improve the scalability properties of Internet routing: address role separation, encapsulation, and mapping. When invented, LISP was targeted at solving the Internet routing scaling problem [RFC4984]. There have now been years of implementations and experiments examining the impact and open questions of using LISP to improve inter-domain routing scalability. Experience has shown that because LISP utilizes mapping and encapsulation technologies, it can be deployed and used for purposes that go beyond routing scalability. For example, LISP provides a mean for a LISP site to precisely control its inter-domain outgoing and incoming traffic, with the possibility to apply different policies to different domains exchanging traffic with it. LISP can also be used to ease the transition from IPv4 to IPv6 as it allows the transport of IPv4 over IPv6 or IPv6 over IPv4. Furthermore, LISP also supports inter-domain multicast.

Leveraging implementation and deployment experience, as well as research work, this document describes, at a high level, the impacts and open questions still seen in LISP. This information is particularly useful for considering future approaches and to support further experimentation to clarify some large open questions (e.g., around the operations). LISP utilizes a tunnel-based data plane and a distributed control plane. LISP requires some new functionalities, such as reachability mechanisms. Because LISP is more than a simple encapsulation technology and is a new technology, until even more deployment experience is gained, some open questions related to LISP deployment and operations remain. As an encapsulation technology, there may be concerns on reduced Maximum Transmission Unit (MTU) size in some deployments. An important impact of LISP is on network operations related to resiliency and troubleshooting. As LISP relies on cached mappings and on encapsulation, resiliency during failures and troubleshooting may be more difficult. Also, the use of encapsulation may make failure detection and recovery slower, and it will require more coordination than with a single, non-encapsulated, routing domain solution.

2. LISP in a Nutshell

LISP relies on three principles: address role separation, encapsulation, and mapping.

The address space is divided into two sets that have different semantic meanings: the Routing Locators (RLOCs) and the Endpoint Identifiers (EIDs). RLOCs are addresses typically assigned from the Provider Aggregatable (PA) address space. The EIDs are attributed to the nodes in the edge networks, by a block of contiguous addresses, which are typically Provider Independent (PI). To limit the scalability problem, LISP only requires the PA routes towards the RLOCs to be announced in the provider infrastructure. Whereas for non-LISP deployments, the EIDs need to be propagated as well.

LISP routers are used at the boundary between the EID and the RLOC spaces. Routers used to exit the EID space (towards the provider domain) are called Ingress Tunnel Routers (ITRs), and those used to enter the EID space (from the provider domain) are called the Egress Tunnel Routers (ETRs). When a host sends a packet to a remote destination, it sends it as in the non-LISP Internet. The packet arrives at the border of its site at an ITR. Because EIDs are not routable on the Internet, the packet is encapsulated with the source address set to the ITR RLOC and the destination address set to the ETR RLOC. The encapsulated packet is then forwarded in the provider domain until it reaches the selected ETR. The ETR de-encapsulates the packet and forwards it to its final destination. The acronym xTR stands for Ingress/Egress Tunnel Router and is used for a router playing these two roles.

The correspondence between EIDs and RLOCs is given by the mappings. When an ITR needs to find ETR RLOCs that serve an EID, it queries a mapping system. With the LISP Canonical Address Format (LCAF) [LISP-LCAF], LISP is not restricted to the Internet protocol for the EID addresses. With LCAF, any address type can be used as EID (the address is only the key for the mapping lookup). LISP can transport, for example, Ethernet frames over the Internet.

An introduction to LISP can be found in [RFC7215]. The LISP specifications are given in [RFC6830], [RFC6833], [LISP-DDT], [RFC6836], [RFC6832], and [RFC6834].

3. LISP for Scaling the Internet Routing Architecture

The original goal of LISP was to improve the scalability properties of the Internet routing architecture. LISP utilizes traffic engineering and stub Autonomous System (AS) prefixes (not announced anymore in the Default-Free Zone (DFZ)), so that routing tables are smaller and more stable (i.e., they experience less churn). Furthermore, at the edge of the network, information necessary to forward packets (i.e., the mappings) is obtained on demand using a pull model (whereas the current Internet BGP model uses a push model). Therefore, the scalability of edge networks is less dependent on the Internet's size and more related to its traffic matrix. This scaling improvement has been proven by several studies (see below). The research studies cited hereafter are based on the following assumptions:

- EID-to-RLOC mappings follow the same prefix size as the current BGP routing infrastructure (current PI addresses only);
- EIDs are used only at the stub ASes, not in the transit ASes; and
- the RLOCs of an EID prefix are deployed at the edge between the stubs owning the EID prefix and the providers, allocating the RLOCs in a PA mode.

The above assumptions are inline with [RFC7215] and current LISP deployments. It is recognized these assumptions may change in the longer term. [KIF13] and [CDLC] explore different EID prefix space sizes and still show results that are consistent and equivalent to the above assumptions.

Quoitin et al. [QIdLB07] show that the separation between locator and identifier roles at the network level improves the routing scalability by reducing the Routing Information Base (RIB) size (up to one order of magnitude) and increases path diversity and thus the traffic engineering capabilities. [IB07] and [KIF13] show, based on real Internet traffic traces, that the number of mapping entries that must be handled by an ITR of a network with up to 20,000 users is limited to few tens of thousands; the signaling traffic (i.e., Map-Request/Map-Reply packets) is in the same order of magnitude similar to DNS request/reply traffic; and the encapsulation overhead, while not negligible, is very limited (in the order of few percentage points of the total traffic volume).

Previous studies consider the case of a timer-based cache eviction policy (i.e., mappings are deleted from the cache upon timeout), while [CDLC] has a more general approach based on the Least Recently Used (LRU) eviction policy, proposing an analytic model for the EID-

to-RLOC cache size when prefix-level traffic has a stationary generating process. The model shows that miss rate can be accurately predicted from the EID-to-RLOC cache size and a small set of easily measurable traffic parameters. The model was validated using four one-day-long packet traces collected at egress points of a campus network and an academic exchange point considering EID prefixes as being of the same size as BGP prefixes. Consequently, operators can provision the EID-to-RLOC cache of their ITRs according to the miss rate they want to achieve for their given traffic.

Results in [CDLC] indicate that for a given target miss ratio, the size of the cache depends only on the parameters of the popularity distribution; the size of the cache is independent of the number of users (the size of the LISP site) and the number of destinations (the size of the EID prefix space). Assuming that the popularity distribution remains constant, this means that as the number of users and the number of destinations grow, the cache size needed to obtain a given miss rate remains constant O(1).

LISP usually populates its EID-to-RLOC cache in a pull mode, which means that mappings are retrieved on demand by the ITR. The main advantage of this mode is that the EID-to-RLOC cache size only depends on the traffic characteristics at the ITR and is independent of the size of the provider domain. This benefit comes at the cost of some delay to transmit the packets that do not hit an entry in the cache (for which a mapping has to be learned). This delay is bound by the time necessary to retrieve the mapping from the mapping system. Moreover, similarly to a push model (e.g., BGP), the pull model induces signaling messages that correspond to the retrieval of mappings upon cache miss. The difference being that the signaling load only depends on the traffic at the ITR and is not triggered by external events such as in BGP. [CDLC] shows that the miss rate is a function of the EID-to-RLOC cache size and traffic generation process, and [CDLC], [SDIB08], and [SDIB08] show from traffic traces that, in practice, the cache miss rate, and thus the signaling rate, remain low.

4. Beyond Scaling the Internet Routing Architecture

LISP is more than just a scalability solution; it is also a tool to provide both incoming and outgoing traffic engineering [S11] [LISP-TE], it can be used as an IPv6 transition at the routing level, and it can be used for inter-domain multicast [RFC6831] [LISP-RE]. Also, LISP has been identified for use to support devices' Internet mobility [LISP-MN] and to support virtual machines' mobility in data centers and multi-tenant VPNs. These last two uses are not discussed further as they are out of the scope of the current LISP Working Group charter.

A key advantage of the LISP architecture is that it facilitates routing in environments where there is little to no correlation between network endpoints and topological location. In service-provider environments, this application is needed in a range of consumer use cases that require an inline anchor to deliver a service to subscribers. Inline anchors provide one of three types of capabilities:

- enable mobility of subscriber endpoints
- enable chaining of middlebox functions and services
- enable functions to be scaled out seamlessly

Without LISP, the approach commonly used by operators is to aggregate service anchors in custom-built boxes. This limits deployments as endpoints can only move on the same mobile gateway, functions can be chained only if traffic traverses the same wire or the same Deep Packet Inspection (DPI) box, and capacity can be scaled out only if traffic fans out to/from a specific load balancer.

With LISP, service providers are able to distribute, virtualize, and instantiate subscriber-service anchors anywhere in the network. Typical use cases for virtualized inline anchors and network functions include Distributed Mobility and Virtualized Evolved Packet Core (vEPC), Virtualized Customer Premise Equipment (vCPE), where functionality previously anchored at a customer premise is now dynamically allocated in the network, Virtualized SGi LAN, Virtual IP Multimedia Subsystems (IMSs), Virtual Session Border Controller (SBC), etc.

ConteXtream [ConteXtream] has been deploying map-assisted overlay networks since 2006, first with a proprietary solution, then evolving to standard LISP. The solution has been deployed in production in three tier-1 operators spanning hundreds of millions of subscribers. Map-assisted overlays had been primarily used to map subscriber flows to services resources dynamically based on profiles and conditions. Specifically, it has been used to map mobile subscribers to value-added/optimization services, broadband subscribers to telephony services, and fixed-mobile subscribers to Broadband Network Gateway (BNG) functions and Internet access services. The LISP map-assisted overlay architecture is used to optimally resolve subscriber to services, functions, instances, and IP overlay aggregation locations on a per-flow basis and just in time.

4.1. Traffic Engineering

In the current (non-LISP) routing infrastructure, addresses used by stub networks are globally routable, and the routing system distributes the routes to reach these stubs. With LISP, the EID prefixes of a LISP site are not routable in the DFZ; mappings are needed in order to determine the list of LISP routers to contact to forward packets. This difference is significant for two reasons. First, packets are not forwarded to a site but to a specific router. Second, a site can control the entry points for its traffic by controlling its mappings.

For traffic engineering purposes, a mapping associates an EID prefix to a list of RLOCs. Each RLOC is annotated with a priority and a weight. When there are several RLOCs, the ITR selects the one with the highest priority and sends the encapsulated packet to this RLOC. If several RLOCs with the highest priority exist, then the traffic is balanced proportionally to their weight among such RLOCs. Traffic engineering in LISP thus allows the mapping owner to have a fine-grained control on the primary and backup path for its incoming and outgoing packet use. In addition, it can share the load among its links. An example of the use of such a feature is described by Saucez et al. [SDIB08], which shows how to use LISP to direct different types of traffic on different links having different capacity.

Traffic engineering in LISP goes one step further, as every Map-Request contains the source EID address of the packet that caused a cache miss and triggered the Map-Request. It is thus possible for a mapping owner to differentiate the answer (Map-Reply) it gives to Map-Requests based on the requester. This functionality is not available today with BGP because a domain cannot control exactly the routes that will be received by domains that are not in the direct neighborhood.

4.2. LISP for IPv6 Co-existence

The LISP encapsulation mechanism is designed to support any combination of address families for locators and identifiers. It is then possible to bind IPv6 EIDs with IPv4 RLOCs and vice versa. This allows transporting IPv6 packets over an IPv4 network (or IPv4 packets over an IPv6 network), making LISP a valuable mechanism to ease the transition to IPv6.

An example is the case of the network infrastructure of a data center being IPv4 only while dual-stack front-end load balancers are used. In this scenario, LISP can be used to provide IPv6 access to servers even though the network and the servers only support IPv4. Assuming

that the data center's ISP offers IPv6 connectivity, the data center only needs to deploy one (or more) xTR(s) at its border with the ISP and one (or more) xTR(s) directly connected to the load balancers. The xTR(s) at the ISP's border tunnels IPv6 packets over IPv4 to the xTR(s) directly attached to the load balancer. The load balancer's xTR de-encapsulates the packets and forwards them to the load balancer, which act as a proxy, translating each IPv6 packet into an IPv4 packet. IPv4 packets are then sent to the appropriate servers. Similarly, when the server's response arrives at the load balancer, the packet is translated back into an IPv6 packet and forwarded to its xTR(s), which in turn will tunnel it back, over the IPv4-only infrastructure, to an xTR connected to the ISP. The packet is then de-encapsulated and forwarded to the ISP natively in IPv6.

4.3. Inter-domain Multicast

LISP has native support for multicast [RFC6831]. From the data-plane perspective, at a multicast-enabled xTR, an EID-sourced multicast packet is encapsulated in another multicast packet and subsequently forwarded in an RLOC-level distribution tree. Therefore, xTRs must participate in both EID and RLOC-level distribution trees. Control-plane wise, since group addresses have no topological significance, they need not be mapped. It is worth noting that, to properly function, LISP-Multicast requires that inter-domain multicast be available.

LISP Replication Engineering (LISP-RE) [LISP-RE] [CDM12] leverages LISP messages [LISP-MULTI-SIGNALING] for multicast state distribution to construct xTR-based inter-domain multicast distribution trees when inter-domain multicast support is not available. Simulations of three different management strategies for low-latency content delivery show that such overlays can support thousands of member xTRs, support hundreds of thousands of end hosts, and deliver content at latencies close to unicast ones [CDM12]. It was also observed that high client churn has a limited impact on performance and management overhead.

Similar to LISP-RE, "Signal-Free LISP Multicast" [LISP-SFM] can be used when the core network does not provide multicast support. But instead of using signaling to build inter-domain multicast trees, signal-free exclusively leverages the map server for multicast state storage and distribution. As a result, the source ITR generally performs head-end replication, but it might also be used to emulate LISP-RE distribution trees.

5. Impact of LISP on Operations and Business Models

Numerous implementation efforts ([IOSNXOS], [OpenLISP], [LISPmob], [LISPClick], [LISPcp], and [LISPfritz]) have been made to assess the specifications, and additionally, interoperability tests [Was09] have been successful. A worldwide large deployment in the international lisp4.net testbed, which is currently composed of nodes running at least three different implementations, will allow us to learn further operational aspects related to LISP.

The following sections distinguish the impact of LISP on LISP sites from the impact on non-LISP sites.

5.1. Impact on Non-LISP Traffic and Sites

LISP has no impact on traffic that has neither LISP origin nor LISP destination. However, LISP can have a significant impact on traffic between a LISP site and a non-LISP site. Traffic between a non-LISP site and a LISP site is subject to the same issues as those observed for LISP-to-LISP traffic but also has issues specific to the transition mechanism that allow the LISP site to exchange packets with a non-LISP site [RFC6832] [RFC7215].

The transition requires setup of proxy tunnel routers (PxTRs). Proxies cause what is referred to as path stretch (i.e., a lengthening of the path compared to the topological shortest path) and make troubleshooting harder. There are still questions related to PxTRs that need to be answered:

- Where to deploy PxTRs? The placement in the topology has an important impact on the path stretch.
- How many PxTRs? The number of PxTRs has a direct impact on the load and the impact of the failure of a PxTR on the traffic.
- What part of the EID space? Will all the PxTRs be proxies for the whole EID space, or will it be segmented between different PxTRs?
- Who operates PxTRs? An important question to answer is related to the entities that will deploy PxTRs: how will they manage their additional Capital Expenditure (CAPEX) / Operating Expenses (OPEX) associated with PxTRs? How will the traffic be carried with respect to security and privacy?

A PxTR will also normally advertise in BGP the EID prefix for which they are proxies. However, if proxies are managed by different entities, they will belong to different ASes. In this case, we need to be sure that this will not cause Multi-Origin AS (MOAS) issues

that could negatively influence routing. Moreover, it is important to ensure that the way EID prefixes will be de-aggregated by the proxies will remain reasonable so as not to contribute to BGP scalability issues.

5.2. Impact on LISP Traffic and Sites

LISP is a protocol based on the map-and-encap paradigm, which has the positive impacts that we have summarized in the above sections. However, LISP also has impacts on operations:

MTU issue: As LISP uses encapsulation, the MTU is reduced; this has implications on potentially all of the traffic. However, in practice, on the lisp4.net network, no major issue due to the MTU has been observed. This is probably due to the fact that current end-host stacks are well designed to deal with the problem of MTU.

Resiliency issue: The advantage of flexibility and control offered by the Locator/ID separation comes at the cost of increasing the complexity of the reachability detection. Indeed, identifiers are not directly routable and have to be mapped to locators, but a locator may be unreachable while others are still reachable. This is an important problem for any tunnel-based solution. In the current Internet, packets are forwarded independently of the border router of the network meaning that, in case of the failure of a border router, another one can be used. With LISP, the destination RLOC specifically designates one particular ETR; hence, if this ETR fails, the traffic is dropped, even though other ETRs are available for the destination site. Another resiliency issue is linked to the fact that mappings are learned on demand. When an ITR fails, all its traffic is redirected to other ITRs that might not have the mappings requested by the redirected traffic. Existing studies [SKI12] [SD12] show, based on measurements and traffic traces, that failure of ITRs and RLOC are infrequent but that when such failure happens, a critical number of packets can be dropped. Unfortunately, the current techniques for LISP resiliency, based on monitoring or probing, are not rapid enough (failure recovery on the order of a few seconds). To tackle this issue, [LISP-PRESERVE] and [LISP-ITR-GRACEFUL] propose techniques based on local failure detection and recovery.

Middleboxes/filters: Because of the increasingly common use of encryption as a response to pervasive monitoring [RFC7258] with LISP providing the option to encrypt traffic between xTRs [LISP-CRYPTO], middleboxes are increasingly likely to be unable to understand encapsulated traffic, which can cause them to drop legitimate packets. In addition, LISP allows triangular or even

rectangular routing, so it is difficult to maintain a correct state even if the middlebox understands LISP. Finally, filtering may also have problems because they may think only one host is generating the traffic (the ITR), as long as it is not de-encapsulated. To deal with LISP encapsulation, LISP-aware firewalls that inspect inner LISP packets are proposed [lispfirewall].

Troubleshooting/debugging: The major issue that LISP experimentation has shown is the difficulty of troubleshooting. When there is a problem in the network, it is hard to pinpoint the reason as the operator only has a partial view of the network. The operator can see what is in its EID-to-RLOC cache/database and can try to obtain what is potentially elsewhere by querying the Map Resolvers, but the knowledge remains partial. On top of that, ICMP packets only carry the first few tens of bytes of the original packet, which means that when an ICMP arrives at the ITR, it might not contain enough information to allow correct troubleshooting. Deployment in the beta network has shown that LISP+ALT [RFC6836] was not easy to maintain and control [CCR13], which explains the migration to LISP-DDT [LISP-DDT], based on a massively distributed and hierarchical approach [CCR13].

Business/operational related: Iannone et al. [IL10] have shown that there are economical incentives to migrate to LISP; however, some questions remain. For example, how will the EIDs be allocated to allow aggregation and hence scalability of the mapping system? Who will operate the mapping system infrastructure and for what benefits? What if several operators run different mapping systems? How will they interoperate or share mapping information?

Reachability: The overhead related to RLOC reachability mechanisms is not known.

6. Security Considerations

A thorough security and threat analysis of LISP is carried out in detail in [RFC7835]. For LISP and other Internet technologies, most of the threats can be mitigated using Best Current Practices, meaning with careful deployment and configuration (e.g., filter), by activating only features that are really necessary in the deployment, and by verifying all the information obtained from third parties. Unless gleaning (Section 6 of [RFC6830] and Section 3.1 of [RFC7835]) features are used, the LISP data plane shows the same level of security as other IP-over-IP technologies. From a security perspective, the control plane remains the critical part of the LISP architecture. To mitigate the threats on the mapping system, authentication should be used for all control-plane messages. The

current specification defines security mechanisms [RFC6836] [LISP-SEC] that can reduce threats in open network environments. The LISP specification defines a generic authentication data field for control-plane messages [RFC6836], which could be used for a general authentication mechanism for the LISP control plane while staying backward compatible.

7. References

7.1. Normative References

[RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, DOI 10.17487/RFC6830, January 2013, <http://www.rfc-editor.org/info/rfc6830>.

[RFC6831] Farinacci, D., Meyer, D., Zwiebel, J., and S. Venaas, "The Locator/ID Separation Protocol (LISP) for Multicast Environments", RFC 6831, DOI 10.17487/RFC6831, January 2013, <http://www.rfc-editor.org/info/rfc6831>.

[RFC6832] Lewis, D., Meyer, D., Farinacci, D., and V. Fuller, "Interworking between Locator/ID Separation Protocol (LISP) and Non-LISP Sites", RFC 6832, DOI 10.17487/RFC6832, January 2013, <http://www.rfc-editor.org/info/rfc6832>.

[RFC6833] Fuller, V. and D. Farinacci, "Locator/ID Separation Protocol (LISP) Map-Server Interface", RFC 6833, DOI 10.17487/RFC6833, January 2013, <http://www.rfc-editor.org/info/rfc6833>.

[RFC6834] Iannone, L., Saucez, D., and O. Bonaventure, "Locator/ID Separation Protocol (LISP) Map-Versioning", RFC 6834, DOI 10.17487/RFC6834, January 2013, <http://www.rfc-editor.org/info/rfc6834>.

[RFC6836] Fuller, V., Farinacci, D., Meyer, D., and D. Lewis, "Locator/ID Separation Protocol Alternative Logical Topology (LISP+ALT)", RFC 6836, DOI 10.17487/RFC6836, January 2013, <http://www.rfc-editor.org/info/rfc6836>.

[RFC7215] Jakab, L., Cabellos-Aparicio, A., Coras, F., Domingo-Pascual, J., and D. Lewis, "Locator/Identifier Separation Protocol (LISP) Network Element Deployment Considerations", RFC 7215, DOI 10.17487/RFC7215, April 2014, <http://www.rfc-editor.org/info/rfc7215>.

[RFC7835] Saucez, D., Iannone, L., and O. Bonaventure, "Locator/ID Separation Protocol (LISP) Threat Analysis", RFC 7835, DOI 10.17487/RFC7835, April 2016, <http://www.rfc-editor.org/info/rfc7835>.

7.2. Informative References

[CCR13] Saucez, D., Iannone, L., and B. Donnet, "A First Measurement Look at the Deployment and Evolution of the Locator/ID Separation Protocol", ACM SIGCOMM Computer Communication Review, Vol. 43, Issue 2, pp. 37-43, DOI 10.1145/2479957.2479963, April 2013.

[CDLC] Coras, F., Domingo, J., Lewis, D., and A. Cabellos, "An Analytical Model for Loc/ID Mappings Caches", IEEE/ACM Transactions on Networking, Vol. 24, Issue 1, pp. 506-516, DOI 10.1109/TNET.2014.2373398, February 2014.

[CDM12] Coras, F., Domingo-Pascual, J., Maino, F., Farinacci, D., and A. Cabellos-Aparicio, "Lcast: Software-defined Inter-Domain Multicast", Computer Networks, Vol. 59, pp. 153-170, DOI 10.1016/j.bjp.2013.10.010, February 2014.

[ConteXtream] ConteXtream Software Company, , "SDN and NFV solutions for carrier networks. (Further details on LISP only through private inquiry.)", <http://www.contextream.com>.

[IB07] Iannone, L. and O. Bonaventure, "On the cost of caching locator/ID mappings", in Proceedings of ACM CoNEXT 2007, DOI 0.1145/1364654.1364663, December 2007.

[IL10] Iannone, L. and T. Leva, "Modeling the economics of Loc/ID Split for the Future Internet", IOS Press, pp. 11-20, DOI 10.3233/978-1-60750-539-6-11, May 2010.

[IOSNXOS] Cisco Systems Inc., "Locator/ID Separation Protocol (LISP)", 2015, <http://lisp4.cisco.com>.

[KIF13] Kim, J., Iannone, L., and A. Feldmann, "Caching Locator/ID mappings: An experimental scalability analysis and its implications", Computer Networks, Vol. 57, Issue 4, DOI 10.1016/j.comnet.2012.11.007, March 2013.

[LISP-CRYPTO] Farinacci, D. and B. Weis, "LISP Data-Plane Confidentiality", Work in Progress, draft-ietf-lisp-crypto-03, September 2015.

[LISP-DDT] Fuller, V., Lewis, D., Ermagan, V., and A. Jain, "LISP Delegated Database Tree", Work in Progress, draft-ietf-lisp-ddt-03, April 2015.

[LISP-ITR-GRACEFUL] Saucez, D., Bonaventure, O., Iannone, L., and C. Filsfils, "LISP ITR Graceful Restart", Work in Progress, draft-saucez-lisp-itr-graceful-03, December 2013.

[LISP-LCAF] Farinacci, D., Meyer, D., and J. Snijders, "LISP Canonical Address Format (LCAF)", Work in Progress, draft-ietf-lisp-lcaf-12, September 2015.

[LISP-MN] Farinacci, D., Lewis, D., Meyer, D., and C. White, "LISP Mobile Node", Work in Progress, draft-meyer-lisp-mn-14, July 2015.

[LISP-MULTI-SIGNALING] Farinacci, D. and M. Napierala, "LISP Control-Plane Multicast Signaling", Work in Progress, draft-farinacci-lisp-mr-signaling-06, February 2015.

[LISP-PRESERVE] Bonaventure, O., Francois, P., and D. Saucez, "Preserving the reachability of LISP ETRs in case of failures", Work in Progress, draft-bonaventure-lisp-preserve-00, July 2009.

[LISP-RE] Coras, F., Cabellos-Aparicio, A., Domingo-Pascual, J., Maino, F., and D. Farinacci, "LISP Replication Engineering", Work in Progress, draft-coras-lisp-re-08, November 2015.

[LISP-SEC] Maino, F., Ermagan, V., Cabellos-Aparicio, A., and D. Saucez, "LISP-Security (LISP-SEC)", Work in Progress, draft-ietf-lisp-sec-10, October 2015.

[LISP-SFM] Moreno, V. and D. Farinacci, "Signal-Free LISP Multicast", Work in Progress, draft-ietf-lisp-signal-free-multicast-01, April 2016.

[LISP-TE] Farinacci, D., Kowal, M., and P. Lahiri, "LISP Traffic Engineering Use-Cases", Work in Progress, draft-farinacci-lisp-te-10, September 2015.

[LISPClick] Saucez, D. and V. Nguyen, "LISP-Click: A Click implementation of the Locator/ID Separation Protocol", 1st Symposium on Click Modular Router, November 2009, <http://hdl.handle.net/2078.1/79067>.

[LISPcp] "LIP6-LISP open source project", 2014, <https://github.com/lip6-lisp>.

[lispfirewall] "LISP and Zone-Based Firewalls Integration and Interoperability", 2014, <http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_data_zbf/configuration/xe-3s/sec-data-zbf-xe-book/sec-zbf-lisp-inner-pac-insp.html>.

[LISPfritz] "Unsere FRITZ!Box-Produkte", 2014, <http://avm.de/produkte/fritzbox/>.

[LISPmob] "An open-source LISP implementation for Linux, Android and OpenWRT", 2015, <http://lispmob.org>.

[OpenLISP] "The OpenLISP Project", 2013, <http://www.openlisp.org>.

[QIdLB07] Quoitin, B., Iannone, L., de Launois, C., and O. Bonaventure, "Evaluating the Benefits of the Locator/Identifier Separation", in Proceedings of MobiArch, Article No. 5, DOI 10.1145/1366919.1366926, August 2007.

[RFC4984] Meyer, D., Ed., Zhang, L., Ed., and K. Fall, Ed., "Report from the IAB Workshop on Routing and Addressing", RFC 4984, DOI 10.17487/RFC4984, September 2007, <http://www.rfc-editor.org/info/rfc4984>.

[RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <http://www.rfc-editor.org/info/rfc7258>.

[S11] Saucez, D., "Mechanisms for Interdomain Traffic Engineering with LISP", PhD Thesis, Universite catholique de Louvain, September 2011, <http://hdl.handle.net/2078.1/92231>.

[SD12] Saucez, D. and B. Donnet, "On the Dynamics of Locators in LISP", in Proceedings of IFIP/TC6 Networking, pp. 385-396, DOI 10.1007/978-3-642-30045-5_29, May 2012.

[SDIB08] Saucez, D., Donnet, B., Iannone, L., and O. Bonaventure, "Interdomain Traffic Engineering in a Locator/Identifier Separation Context", in Proceedings of Internet Network Management Workshop, DOI 10.1109/INETMW.2008.4660330, October 2008.

[SKI12] Saucez, D., Kim, J., Iannone, L., Bonaventure, O., and C. Filsfils, "A Local Approach to Fast Failure Recovery of LISP Ingress Tunnel Routers", in Proceedings of IFIP Networking 2012, pp. 397-408, DOI 10.1007/978-3-642-30045-5_30, May 2012.

[Was09] Wasserman, M., "LISP Interoperability Testing", IETF 76, LISP WG Presentation, November 2009.

Acknowledgments

Thanks to Deborah Brungard, Ben Campbell, Spencer Dawkins, Stephen Farrel, Wassim Haddad, Kathleen Moriarty, and Hilarie Orman for their thorough reviews, comments, and suggestions.

The people that contributed to this document are Alia Atlas, Sharon Barkai, Ron Bonica, Ross Callon, Vince Fuller, Joel Halpern, Terry Manderson, and Gregg Schudel.

The work of Luigi Iannone has been partially supported by the ANR 13 INFR 0009 LISP-Lab Project <http://www.lisp-lab.org>.

Authors' Addresses

Damien Saucez
INRIA
2004 route des Lucioles BP 93
06902 Sophia Antipolis Cedex
France

Email: damien.saucez@inria.fr

Luigi Iannone
Telecom ParisTech
23, Avenue d'Italie, CS 51327
75214 Paris Cedex 13
France

Email: ggx@gigix.net

Albert Cabellos
Technical University of Catalonia
C/Jordi Girona, s/n
08034 Barcelona
Spain

Email: acabello@ac.upc.edu

Florin Coras
Technical University of Catalonia
C/Jordi Girona, s/n
08034 Barcelona
Spain

Email: fcoras@ac.upc.edu