Network Working Group
Internet-Draft
Intended status: Informational
Expires: July 12, 2015

D. Saucez
INRIA
L. Iannone
Telecom ParisTech
A. Cabellos
F. Coras
Technical University of Catalonia
January 8, 2015

LISP Impact
draft-ietf-lisp-impact-00.txt

Abstract

The Locator/Identifier Separation Protocol (LISP) aims at improving the Internet scalability properties leveraging on three simple principles: address role separation, encapsulation, and mapping. In this document, based on implementation, deployment, and theoretical studies, we discuss the impact that deployment of LISP can have on both the Internet in general and for the end-users in particular.

Table of Contents

1. Introduction

The Locator/Identifier Separation Protocol (LISP) relies on three simple principles to scale the Internet: address role separation, encapsulation, and mapping. The main goal of LISP is to make the Internet more scalable by reducing the number of prefixes announced in the Default Free Zone (DFZ) as well as its related churn. As LISP relies on mapping and encapsulation, it turns out that it provides more benefits than just 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 the different domains exchanging traffic with it. LISP can also be used to ease the transition from IPv4 to IPv6 as it allows to transport IPv4 over IPv6 or IPv6 over IPv4. Furthermore, LISP also provides a solution to perform inter-domain multicast.

This document discusses the impact of LISP's deployment on the Internet and on end-users and shows the consequences of the interworking infrastructure in path stretch. There still are many, economical rather than technical, open questions related to the deployment of such infrastructure. Moreover, encapsulation may raise some issues (that do not have a real impact in practice) because it

reduces the Maximum Transmission Unit (MTU) size. An important impact of LISP on network operations is related to resiliency and troubleshooting. Indeed, as LISP relies on cached mappings and on encapsulation, troubleshooting is harder than in the traditional Internet. Also, end-to-end encapsulation stresses resiliency as it makes failure detection and recovery slower than with hop-by-hop routing.

2. LISP in a nutshell

The Locator/Identifier Separation Protocol (LISP) relies on three simple principles: address role separation, encapsulation, and mapping.

Semantics of address are separated in two: the Routing Locators (RLOCs) and the Endpoint Identifiers (EIDs). RLOCs are assigned from the address space of the Internet service providers (PA). The EIDs are attributed, to the nodes in the edge network, by block of contiguous addresses extracted from the EID Space. To limit the scalability problem of today's Internet, only the routes towards the RLOCs are announced in the Internet while EIDs are also propagated today.

LISP routers are used at the boundary between the EID and the RLOC spaces. Routers used to exit the EID space are called Ingress Tunnel Router (ITRs) and those used to enter the EID space the Egress Tunnel Routers (ETRs). When a host sends a packet to a remote destination, it sends it as in today's Internet. The packet eventually 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 Internet until it reaches the selected ETR. The ETR decapsulates the packet and forwards it to its final destination. The acronym xTR for Ingress/ Egress tunnel router 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 the mapping system. It is worth noticing that with the LISP Canonical Address Format (LCAF) [I-D.ietf-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 the key for the mapping lookup) and LISP can then transport, for example, Ethernet frames over the Internet.

A more thorough introduction to LISP can be found in [I-D.ietf-lisp-introduction]. The complete specifications are given

in [RFC6830], [RFC6833], [I-D.fuller-lisp-ddt], [RFC6836], [RFC6832], [RFC6834], and [I-D.ietf-lisp-sec].

3. LISP for scaling the Internet

The first goal of LISP is to scale the Internet. LISP improves the Internet's scalability because traffic engineering and stub AS prefixes are not propagated in the DFZ, so routing tables are smaller and more stable (i.e., less affected by churn). Also, at the edge network, information necessary to forward packets (i.e., the mappings) is usually obtained on demand using a pull model. Therefore, for each edge network they scale with the traffic matrix of the edge network and are independent of the Internet's size. This scaling improvement is proven by several works.

Quoitin et al. show in [QIdLB07] that the separation between locator and identifier roles at the network level improves the routing scalability by reducing the RIB size (up to one order of magnitude) and increases the path diversity and thus the traffic engineering capabilities. In addition, Iannone and Bonaventure show in [IB07] that the number of mapping entries that must be supported at an ITR of a 10,000 users campus network is limited and does not represent more that 3 to 4 Megabytes of memory. Furthermore, they show that signaling traffic (i.e., Map-Request/Map-Reply packets) is in the same order of magnitude like DNS requests traffic and that encapsulation overhead, while not negligible, is very limited (in the order of few percentage points of the total traffic volume). Similarly, Kim et al. show that the EID-to-RLOC cache size should not exceed 14 MB for an ITR responsible of more than 20,000 residential ADSL users at a large ISP [KIF11]. [IB07], [KIF11] rely on BGP and traffic traces to determine the number of entries to keep in the EID-to-RLOC cache. In both papers, the size of the cache is inferred from the number of entries by considering that every EID is associated with two or three locators. [S11] confirms these results by looking at the distribution of the number of locators per EID if LISP were deployed in the 2010's Internet. The assumptions in these studies are:

- contiguous addresses tend to be used similarly, EID prefixes follow the current BGP prefixes decomposition;
- EIDs are used only at the stub ASes, not in the transit ASes;
- the RLOCs of an EID prefix are deployed at the edge between the stubs owning the EID prefix and the providers and locator addresses are allocated in a Provider Aggregetable (PA) mode.

While all previous studies consider the case of a timer-based cache eviction policy (i.e., mappings are deleted from the cache upon timeout), [CCD12] generalizes the caching discussion for the Least Recently Used (LRU) eviction policy and proposes 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 BGP-prefix granularity. 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.

The results indicate that for a given miss ratio, cache size only depends on the parameters of the popularity distribution and is in fact 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).

Under normal user traffic, miss-ratio decreases at an accelerated pace with cache size and finally settles to a power-law decrease. However, [CDLC] extends the model to account for scanning attacks, whereby attackers generate a constant flux of packets according to random scans of the destination prefix space and shows that miss-ratios are be very high and independent of cache size. In fact, if the attack is merely 1% of the legitimate traffic, the miss rate does not drop under 1% as long as the cache cannot accommodate the whole prefix space. Locality measurements also suggested that LRU eviction policy should be close to optimal.

TBD: add a paragraph to explain thhe operational difference while dealing with a pull model instead of a push.

4. Beyond scaling the Internet

Even though it is its main goal, LISP is more than just a scalability solution, it is also a tool to provide both incoming and outgoing traffic engineering [S11], can be used as an IPv6 transition at the routing level, and for inter-domain multicast [RFC6831], [I-D.coras-lisp-re]. LISP has also proven to be a good protocol for mobility of devices in the Internet [I-D.meyer-lisp-mn] or even virtual machine mobility in data centers and multi-tenant VPN, however, we don't further discuss in details the two last points as they are out of the scope of the charter.

Lisp architecture facilitates routing in environments where there is little to no correlation between network endpoints and topological location. In service provider environment this use is evident in a range of consumer use cases which require an inline anchor in-order to deliver a service to a subscribers. Inline anchors provide one of three types of capabilities:

- enable mobility of subscriber end points
- enable chaining of middle-box functions
- enable seamless scale-out of functions

Without LISP operators are forced to centralize service anchors in custom built special boxes. This means that end-points can move as long as their traffic ends up on the same mobile gateway, functions can be chained as long as all traffic traverses the same wire or the same DPI box, and capacity can scale out as long as traffic fans out to and form a specific load balancer.

With LISP service providers are able to distribute, virtualize, and insatiate subscriber-service anchors anywhere in the network. Typical use cases that Virtualize inline anchors and network functions include: Distributed Mobility and Virtualized Evolved Packet Core (vEPC), where centralization makes way to distributed and virtualized inline anchoring of mobility, Virtualized Customer Premise Equipment or vCPE, where functionality previously anchored at customer prem is now dynamically allocated in-network, Virtualized SGi LAN, where value added mobile services previously anchored inside full-stack boxes or anchored to physical wires with permutation setups aka "Rails", Virtual IMS and Virtual SBC, etc.

Current deployments by ConteXtream, using a pre standards (designed 2006) based architecture, support a total of 100 millions subscribers with such an architecture. A deployment at a tier-1 US Mobile operator over 50 millions subscribers provides a 39% download rate improvement over LTE.

4.1. Traffic engineering

In today's Internet, stub networks are globally routable and the routing system distributes the routes to reach these stubs. On the contrary, the EID prefixes of a LISP site are not routable on the Internet and mappings are needed to determine the list of LISP routers to contact to send them packets. The difference is significant for two reasons. First, packets are not sent to a site but to a specific ingress router. Second, a site can control the entry points for its traffic by controlling its mappings.

For traffic engineering purpose, 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 lowest priority value and sends the encapsulated packet to this RLOC. If several such RLOCs exist, then the traffic is balanced proportionally to their weight among the RLOCs with the lowest priority value. Traffic engineering in LISP thus allows the mapping owner to have a fine-grained control on the primary and backup path its incoming and outgoing packets use. In addition, it can share the load among its links. An example of the use of such a feature is described in [SDIB08], where Saucez et al. show 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 locators and identifiers address family. 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.

A not so uncommon example is the case of the network infrastructure of a datacenter 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 datacenter's ISP offers IPv6 connectivity, the datacenter 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 decapsulates the packets and forward them to the load balancer, which act as proxies, translating each IPv6 packet into an IPv4. 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 decapsulated 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 a 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 inter-domain, LISP-Multicast requires that inter-domain multicast be prior deployed.

[I-D.coras-lisp-re] and [CDM12] propose a technique to construct xTR based inter-domain multicast distribution trees. Simulations of three different management strategies for low latency content delivery show that such overlays can support thousands of member xTRs, 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.

5. Impact of LISP on operations and business model

Important implementation efforts ([IOSNXOS], [OpenLISP], [LISPmob], [LISPClick], [LISPcp], and [LISPfritz]) have been made to assess the specifications and interoperability tests [Was09] have been a success. World-wide large deployment in the international lisp4.net testbed, which is currently composed of nodes running at least three different implementations, allows to learn operational matters related to LISP.

We have to 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 which 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 are subject to the same issues than those observed for LISP-to-LISP traffic (cf infra) but also have issues specific to the transition mechanism that allow LISP site to exchange packets with non-LISP site ([RFC6832], [I-D.ietf-lisp-deployment]).

Indeed, the transition requires to setup proxy tunnel routers (PxTRs). PxTRs do not cause particular technical issue. However, by definition proxies cause path stretch and make troubleshooting harder. There are still big questions related to PxTRs that have 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 PxTR 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 to operate PxTRs? The IETF does not aim at providing business model hints, however, an important question to answer is related to the entities that will deploy PxTRs, how they will manage their CAPEX/OPEX and how the traffic will be carried with respect for the security and privacy.

PxTR also normally have to advertise in BGP the EID prefix they are proxy for. However, if proxies are managed by different entities, they will belong to different ASes. In this case, we have to be sure that it will not cause MOA issues that could negatively influence routing. Moreover, we have to be sure that the way EID prefixes will be deaggregated by the proxies will remain reasonable to not take part in the 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 effects that we have given in the sections above. However, by design, LISP also has side impact on operations:

MTU issue: as LISP uses encapsulation, the MTU is reduced, this has implication on potentially all 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 designate 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 yet the mappings for the redirected traffic. The study in [SKI12] and [SD12] show, based on measurements and traffic traces, that failure of ITRs and RLOC are infrequent but that when such failure happens, an important number of packet can be dropped. Unfortunately, the current techniques for LISP resiliency, based on monitoring or probing are not rapid enough (failure recovery of the order of a few seconds). To tackle this issue [I-D.bonaventure-lisp-preserve] and [I-D.saucez-lisp-itr-graceful] propose techniques based on local failure detection and recovery.

Middle boxes/filters: because of encapsulation, the middle boxes might not understand the traffic which can cause firewall to drop legitimate packets. In addition, LISP allows triangular or even rectangular routing, so it is hard to maintain a correct state even if the middle box perfectly understands LISP. Finally, filtering might also have problems because they might think only one host is generating the traffic (the ITR), as long as it is not decapsulated. To deal with LISP encapsulation, LISP aware firewalls that inspect inner LISP packets are proposed [lispfirewall].

Troubleshooting/debugging: the major issue years of LISP experimentation have shown is the difficulty of troubleshooting. When there is a problem in the network, it is hard to pin-point 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 is too small, which means that when an ICMP arrives at the ITR, it might not contain enough information to make correct troubleshooting. Interestingly, deployment in the beta network has shown that LISP+ALT was not easy to maintain and control, which explains the migration to LISP-DDT [I-D.fuller-lisp-ddt].

Business: the IETF is not aiming at providing business models. However, even though [IL10] shown that there is economical incentives to migrate to LISP, some questions are on hold. 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 benefit?

6. IANA Considerations

This document makes no request to the IANA.

7. Security Considerations

Security and threats analysis of the LISP protocol is out of the scope of the present document. A thorough analysis of LISP security threats is detailed in [I-D.ietf-lisp-threats].

8. Acknowledgments

The people that contributed to this document are Sharon Barkai, Vince Fuller, Joel Halpern, Terry Manderson, and Gregg Schudel.

9. References

9.1. Normative References

[I-D.fuller-lisp-ddt]
Fuller, V., Lewis, D., Ermagan, V., and A. Jain, "LISP Delegated Database Tree", draft-fuller-lisp-ddt-04 (work in progress), September 2012.

[I-D.ietf-lisp-deployment]
Jakab, L., Cabellos-Aparicio, A., Coras, F., Domingo-Pascual, J., and D. Lewis, "LISP Network Element Deployment Considerations", draft-ietf-lisp-deployment-12 (work in progress), January 2014.

[I-D.ietf-lisp-sec]
Maino, F., Ermagan, V., Cabellos-Aparicio, A., and D. Saucez, "LISP-Security (LISP-SEC)", draft-ietf-lisp-sec-07 (work in progress), October 2014.

[RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, January 2013.

[RFC6831] Farinacci, D., Meyer, D., Zwiebel, J., and S. Venaas, "The Locator/ID Separation Protocol (LISP) for Multicast Environments", RFC 6831, January 2013.

[RFC6832] Lewis, D., Meyer, D., Farinacci, D., and V. Fuller, "Interworking between Locator/ID Separation Protocol (LISP) and Non-LISP Sites", RFC 6832, January 2013.

[RFC6833] Fuller, V. and D. Farinacci, "Locator/ID Separation Protocol (LISP) Map-Server Interface", RFC 6833, January 2013.

[RFC6834] Iannone, L., Saucez, D., and O. Bonaventure, "Locator/ID Separation Protocol (LISP) Map-Versioning", RFC 6834, January 2013.

[RFC6836] Fuller, V., Farinacci, D., Meyer, D., and D. Lewis, "Locator/ID Separation Protocol Alternative Logical Topology (LISP+ALT)", RFC 6836, January 2013.

9.2. Informative References

[CCD12] Coras, F., Cabellos-Aparicio, A., and J. Domingo-Pascual, "An Analytical Model for the LISP Cache Size", In Proc. IFIP Networking 2012, May 2012.

[CDLC] Coras, F., Domingo, J., Lewis, D., and A. Cabellos, "An Analytical Model for Loc/ID Mappings Caches", Technical Report http://arxiv.org/pdf/1312.1378v2.pdf, 2013.

[CDM12] Coras, F., Domingo-Pascual, J., Maino, F., Farinacci, D., and A. Cabellos-Aparicio, "Lcast: Software-defined Inter-Domain Multicast", Technical Report, Universitat Politecnica de Catalunya, 2012, July 2012.

[I-D.bonaventure-lisp-preserve] Bonaventure, O., Francois, P., and D. Saucez, "Preserving the reachability of LISP ETRs in case of failures", draft-bonaventure-lisp-preserve-00 (work in progress), July 2009.

[I-D.chiappa-lisp-architecture] Art, Y., "An Architectural Perspective on the LISP Location-Identity Separation System", draft-chiappa-lisp-architecture-01 (work in progress), July 2012.

[I-D.coras-lisp-re] Coras, F., Cabellos-Aparicio, A., Domingo-Pascual, J., Maino, F., and D. Farinacci, "LISP Replication Engineering", draft-coras-lisp-re-06 (work in progress), October 2014.

[I-D.ietf-lisp-introduction]
Cabellos-Aparicio, A. and D. Saucez, "An Architectural Introduction to the Locator/ID Separation Protocol (LISP)", draft-ietf-lisp-introduction-09 (work in progress), November 2014.

[I-D.ietf-lisp-lcaf]
Farinacci, D., Meyer, D., and J. Snijders, "LISP Canonical Address Format (LCAF)", draft-ietf-lisp-lcaf-07 (work in progress), December 2014.

[I-D.ietf-lisp-threats]
Saucez, D., Iannone, L., and O. Bonaventure, "LISP Threats Analysis", draft-ietf-lisp-threats-11 (work in progress), December 2014.

[I-D.meyer-lisp-mn]
Farinacci, D., Lewis, D., Meyer, D., and C. White, "LISP Mobile Node", draft-meyer-lisp-mn-11 (work in progress), July 2014.

[I-D.saucez-lisp-itr-graceful]
Saucez, D., Bonaventure, O., Iannone, L., and C. Filsfils, "LISP ITR Graceful Restart", draft-saucez-lisp-itr-graceful-03 (work in progress), December 2013.

[IB07] Iannone, L. and O. Bonaventure, "On the cost of caching locator/id mappings", In Proc. ACM CoNEXT 2007, December 2007.

[IL10] Iannone, L. and T. Leva, "Modeling the economics of Loc/ID Separation for the Future Internet", Book Chapter, Towards the Future Internet - Emerging Trends from the European Research, IOS Press, May 2010.

[IOSNXOS] Cisco Systems Inc., , "Locator/ID Separation Protocol (LISP)", http://lisp4.cisco.com, 2013.

[KIF11] Kim, J., Iannone, L., and A. Feldmann, "Deep dive into the lisp cache and what isps should know about it", In Proc. IFIP Networking 2011, May 2011.

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

[LISPcp] "The lip6-lisp Project", https://github.com/lip6-lisp/, 2014.

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

[LISPmob] "LISP Mobile Node for Linux", http://lispmob.org, 2013.

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

[QIdLB07] Quoitin, B., Iannone, L., de Launois, C., and O. Bonaventure, "Evaluating the benefits of the locator/identifier separation", In Proc. ACM MobiArch 2007, May 2007.

[S11] Saucez, D., "Mechanisms for Interdomain Traffic Engineering with LISP", PhD Thesis, Universite catholique de Louvain, 2011, October 2011.

[SD12] Saucez, D. and B. Donnet, "On the Dynamics of Locators in LISP", In Proc. IFIP Networking 2012, May 2012.

[SDIB08] Saucez, D., Donnet, B., Iannone, L., and O. Bonaventure, "Interdomain Traffic Engineering in a Locator/Identifier Separation Context", In Proc. of Internet Network Management Workshop, 2008, 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 Proc. IFIP Networking 2012, May 2012.

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

[lispfirewall] "LISP and Zone-Based Firewalls Integration and Interoperability", 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, 2014.

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: luigi.iannone@telecom-paristech.fr

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

Email: fcoras@ac.upc.edu

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

Email: fcoras@ac.upc.edu