

LISP-MSX: Decentralized Interconnection of Independent LISP Mapping Systems

Mohamed Boucadair, Christian Jacquenet, Dung Chi Phung, Stefano Secci

► **To cite this version:**

Mohamed Boucadair, Christian Jacquenet, Dung Chi Phung, Stefano Secci. LISP-MSX: Decentralized Interconnection of Independent LISP Mapping Systems. IEEE Communications Magazine, Institute of Electrical and Electronics Engineers, 2019, 57 (1), pp.35-41. 10.1109/MCOM.2018.1701323 . hal-01989714

HAL Id: hal-01989714

<https://hal.archives-ouvertes.fr/hal-01989714>

Submitted on 4 Jan 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

LISP-MSX: Decentralized Interconnection of Independent LISP Mapping Systems

Abstract—We present in this paper a novel solution for the interconnection of LISP (Locator/Identifier Separation Protocol) mapping systems. Our solution, named LISP-MSX, differs from existing approaches in that it allows for complete mapping systems technology independence and for their decentralized interconnection, by means of novel control-plane primitives to LISP and routing protocols, hence guaranteeing faster mappings resolution

I. INTRODUCTION

The Internet growth can be assessed by the size of the routing and forwarding tables maintained by the routers that keep a global, topological view of the Internet, i.e., the whole set of IP routes that can reach any terminal connected to the Internet. Such routers compose the Default-Free Zone (DFZ) of the Internet. The aforementioned growth has evolved exponentially for many years¹: there were approximately 10,000 IPv4 routes in 1994 and there are now more than 700,000 of such routes. Likewise, there were a few hundreds of IPv6 routes before 2004 and there are now more than 52,000 routes [<http://cidr-report.org>]

Among the various proposals discussed over the years to improve Internet traffic forwarding efficiency, those that consist in separating the information that is specific to the location where a terminal is connected to the Internet (“where”) from the information that is specific to the identity of the terminal (“who”) have attracted a growing interest within the Internet community. It is generally admitted that the ability to separate the “where” from the “who” allows to get rid of a single address space suffering from prefix de-aggregation, a phenomenon behind the routing table size increase. Multiple Identifier/Locator split addressing protocols were discussed in the last two decades, as documented in [LISPsurvey]. Among them, the Locator/ID Separation Protocol (LISP) differentiates from most of the other approaches in that it does not imply any modification of terminal devices. As few other protocols, LISP underwent standardization for several years [lisp].

The large majority of the Identifier/Locator split protocols need a Mapping System that maintains mappings between the Identifier and the Locator information, and provides mapping resolution services accordingly [MPsurvey]. Several LISP mapping database systems were proposed, but the Delegated Database Tree (LISP-DDT) [lispDDT] is the one currently deployed by operational implementations -

e.g., the Cisco IOS. LISP-DDT proposes a hierarchical resolution model like the DNS (Domain Name Service) system. Such a hierarchical structure may affect resolution times, besides raising political concerns due to potential country-centric management (e.g., DNS-like governance), where the mastering of root servers can influence the quality of the resolution service at the Internet scale. In LISP-DDT, when a mapping resolution request is issued, it is relayed from a resolver node to another one, passing through a DDT, until it reaches an authoritative server. Alternative proposals were discussed, such as ALT (Alternative LISP Topology) [lispALT], which however mandates a parallel node-disjoint separation for the control-plane, with distinct BGP (Border Gateway Protocol) routers.

In this article, we propose a mapping system interconnection infrastructure, named LISP-MSX, which can interoperate transparently with LISP-DDT, while being capable to super-setting it to more directly interconnecting independent mapping system without following a resolution hierarchy. More precisely, we specify and experiment novel LISP control-plane primitives and route discovery protocol extensions in support of (i) a decentralization of the mapping system interconnection and resolution process, (ii) decreasing the mapping resolution latency, and (iii) better scaling of the overall mapping system. About the first aspect, the motivation to adopt LISP-MSX is to avoid falling in a similar situation than the DNS one, with three countries controlling the whole DNS root domain, with many operators willing to get freed by such a dependency. LISP-MSX supports for decentralized governance of mapping system, allowing for customized mapping system implementations within provider boundaries. The framework is based upon new mechanisms to dynamically select remote LISP Mapping Systems, negotiate and establish interconnect agreements with them, and optimize connectivity service operation.

II. LISP BACKGROUND AND OPERATIONS

A. LISP REMAINDER

LISP operation relies upon the manipulation of two identifiers: the Routing LOCator (RLOC), assigned to network topology attachment points, and the Endpoint IDentifiers (EID), assigned to terminal devices independently of the network topology. LISP forwarding uses mapping functions that associate EIDs with RLOCs as well as an encapsulation scheme. As such, LISP does not mandate any specific modification of EID terminals: they use legacy

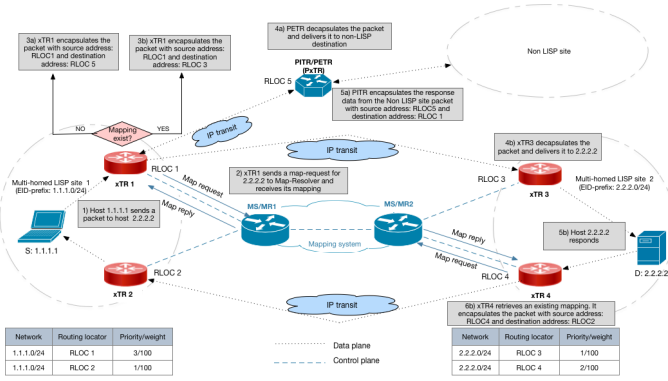


Fig. 1. Example of LISP communications between two LISP sites.

IP addressing and forwarding. Routers that compose a LISP network are called Tunnel Routers. They are responsible for encapsulating/de-capsulating LISP packets. A LISP packet includes the source and destination IP addresses of the RLOCs. Forwarding decisions in the LISP network are made according to the RLOC information, as shown in Fig. 1. A specific transport layer is used to identify the LISP port number and a LISP shim header is used to carry management information.

A Tunnel Router typically connects a LISP site to the network, and depending on the traffic direction, it behaves as an Ingress Tunnel Router (ITR) or an Egress Tunnel Router (ETR). The term xTR is used to refer to a generic TR role (ITR/ETR).

Map-Resolvers (MR), Map-Servers (MS), and other components like authorization and subscription servers are part of a Mapping System. While a MS learns about authoritative EID-to-RLOC mappings from ETRs by means of Map-Register messages [lisp], and records them in the mapping database, a MR processes LISP Map-Requests [lisp] sent by ITRs and solicits MS accordingly to resolve EID-to-RLOC mappings. The mapping resolution service provided by the Mapping System therefore helps xTRs to populate and update their mapping tables.

In order to allow for global reachability, Proxy Ingress/Egress Tunnel Routers (PxTR) [lispPxtr] are deployed to handle traffic between non-LISP and LISP sites. As such, a Proxy Ingress Tunnel Router behaves as an ITR on behalf of non-LISP sites that send packets to destinations located in LISP sites. Likewise, a Proxy ETR behaves as an ETR on behalf of LISP sites that send traffic to destinations located in non-LISP sites.

As shown in Fig. 1, hosts “S” and “D” are assigned an address extracted from the corresponding site’s EID-prefixes (which does not need to be injected in the DFZ). These EID prefixes are registered into the LISP mapping system. For example, the host S in LISP site 1 (EID 1.1.1.1) has to communicate with the host D (EID 2.2.2.2) in LISP site 2. It sends normal IP packets with source and destination IP addresses set to 1.1.1.1 and 2.2.2.2, which reach xTR1 (acting as the ITR of LISP sites). Upon receiving the first packet, xTR1

checks its EID/RLOC mapping table to make its forwarding decision; if no entry is found, it solicits the LISP mapping system to retrieve the Routing Locators (RLOCs) of the destination: by sending a Map-Request message to the Map-Resolver and getting the Map-Reply message. If the Map-Reply contains a positive mapping record, the packet is encapsulated by that ITR and forwarded towards an RLOC of an ETR of D (the destination RLOC of “D” is chosen based on traffic engineering metrics associated to the mapping, a priority and a weight for each RLOC; the lowest value of the priority wins, and in case of many equal values, load-balancing is done accordingly to the weight metrics). The ETR then decapsulates the packet and forwards it natively to D. Note that, if no EID-to-RLOC mapping is available at the Mapping System, there are two possibilities: the first possibility is that the packet is encapsulated by that ITR and forwarded towards a PETR, if a PETR is set, where it is decapsulated and forwarded natively. In the second case, if a PETR is not set, the traffic is forwarded natively, assuming that the destination EID is reachable via legacy IP routing.

B. CHALLENGES OF LISP OPERATION AT THE INTERNET SCALE

The deployment of LISP networks at the scale of the Internet raises several issues that may affect the overall quality of a LISP connectivity service. Various LISP players (network operators, service providers, etc.) are likely to deploy and operate different LISP Mapping Systems [lispService]. Indeed, many proposals were investigated for the past few years, including mobile core networks [evolPacket], software-defined networks [lispSDN], and prefix de-aggregation control practices [interdomainLatin], leading to independent Mapping Systems that may benefit from interconnecting with each other.

Furthermore, Multiple Mapping Systems will co-exist for other reasons, e.g., to avoid country-centric governance, allow for various technologies to implement the mapping service, take advantage of new business opportunities encourage service innovation, etc. The lack of clear policies for the management and operation of the LISP Mapping Systems may encourage such practices.

Moreover, because the LISP Mapping System may provide service differentiation opportunities, IP access and transit providers may be tempted to operate a (local) Mapping System. Mapping Service Providers may offer advanced services to their customers such as the maintenance of local caches, or the update of ITR mapping entries that match some criteria requested by a LISP-enabled network. MS providers may also ensure that mapping resolution requests are redirected to the closest Map-Resolvers, whereas the structuring of the mapping resolution service is meant to optimize mapping resolution times, avoid the loss of the first packet, etc.




Fig. 2. Example of LISP-MSX interconnection between two Mapping Systems.

As represented in Fig. 2, a LISP Mapping System may handle one or multiple prefixes that belong to one or multiple Autonomous Systems (ASes). Distinct flavours of Mapping Systems may be deployed; each may rely upon specific technology. As such, a clear interface to ease interconnection between these realms is needed.

A hierarchy in the Mapping System organization for business, governance, control, and regulatory purposes, in particular, is likely. In such contexts, a Mapping System may maintain (a portion of) a global mapping table. An efficient and scalable LISP deployment within an inter-domain context for traffic engineering purposes heavily relies upon the availability of an inter-domain Mapping System that spans several domains. From this perspective, the success of a global LISP adoption and deployment will mainly depend on how LISP-enabled domains (e.g., an Autonomous System or a simple local area network) will graft to existing Mapping Systems that can guarantee a global reachability scope. To minimize the risk of a fragmented Mapping System that would jeopardize the overall efficiency of an inter-domain LISP routing system, there is a need to encourage and facilitate the coordination of participating Mapping Systems.

III. A FRAMEWORK FOR IMPROVING LISP OPERATION AT LARGE SCALE

Each time there is a need to interconnect two infrastructures owned and managed by distinct entities, a process offering negotiation, interconnection and invocation features is desirable. This process can be static (e.g., the current practice for AS interconnection), but a more dynamic approach would be valuable for the sake of highly automated services and delivery. We propose in the following a framework in this direction for LISP Mapping Systems Interconnection, we refer to as LISP-MSX.

A. AN INTERCONNECT FRAMEWORK FOR A GLOBAL MAPPING SYSTEM

In order to extend the reachability of LISP EIDs beyond the boundaries of a single Mapping System, we aim at proposing a framework that does not require to change xTR behaviour such that an xTR would query multiple Mapping Systems

concurrently (i.e., configured with multiple mapping servers of independent Mapping Systems). These Mapping Systems need to interconnect to extend the reachability scope and avoid pressure on PxTR devices. Also, various Mapping Systems encourage the enforcement of policies that aim at optimizing LISP forwarding: for example, policies that consist in avoiding the solicitation of specific domains or regions (e.g., for security reasons).

It is essential to encourage the deployment and the operation of a global Mapping System at the scale of the Internet instead of a fragmented Mapping System. Fig. 2 depicts a LISP-MSX scenario: while domains 1 and 2 use Mapping System 1, domain 4 uses Mapping System 2. Mapping Systems 1 and 2 are independent, meaning that the LISP traffic exchanged between node N1 and node N2 should use the PxTR. By interconnecting both Mapping Systems, communications between N1 and N2 can be natively LISP-forwarded without invoking any PxTR. Moreover, optimizing such LISP interconnection can also reduce the mapping resolution time compared to the use of a centralized, hierarchical Mapping System such as LISP-DDT.

B. LISP-MSX FUNCTIONAL BLOCKS

The settlement of LISP Mapping System interconnects is decomposed into several functional blocks, as represented in Fig. 3:

- **Discovery and Advertisement:** Discover and Advertise LISP Mapping Systems that are willing to interconnect as well as those that are ready to service LISP-enabled networks. A leaf LISP-enabled network may subscribe to the mapping service provided by one or several Mapping Service Providers. In Fig. 3, Mapping System 2 advertises its reachability information to Mapping System 1.
- **Negotiation:** We identify the mapping negotiation as a viable approach, as it allows getting rid of the need of manual configurations as is the case for the current LISP specification, in particular for the configuration of MSs and DDT roots [**bootstrapping**]. The goal of the Negotiation block is to negotiate interconnection agreements with remote Mapping Service Providers. The same mechanism can be used by a LISP-enabled network to subscribe to one or multiple Mapping Systems. Subscribing to multiple Mapping Systems is meant to enhance the robustness of the connectivity service. The contribution of each player involved in the provisioning and the operation of a LISP-based connectivity forwarding service needs to be rationalized so that clear interfaces are defined and adequate mechanisms for troubleshooting, diagnosis and repair purposes can be easily implemented and adopted. The inability of identifying what is at the origin of the degradation of a LISP connectivity service is seen as one of the hurdles that are likely to jeopardize LISP deployments at the scale of the Internet. The interconnection agreement can be unidirectional

or bi-directional. Dedicated technical clauses may be included in the interconnect agreements to specify whether advanced primitives (such as bulk mapping transfer or record notifications) are supported. Also, the agreement specifies how requests are rate-limited.

- **Mapping System Interconnect:** Implements interconnect agreements with remote Mapping Systems to facilitate the exchange of mapping records between Mapping Systems. The entries of the mapping tables (or a part thereof) are exchanged between these Mapping Systems so that Map-Request messages can be processed as close to the LISP leaf networks as possible.
- **Service Invocation:** Invoke a peer Mapping System for mapping records resolution, in particular. Other services can be offered by the Mapping System, e.g., assist the forwarding of the first packet before a mapping entry is available in the xTR cache.




Fig. 3. Representation of four functional blocks for LISP-MSX operations: Discovery, Negotiation, Interconnection, Invocation).

Also, the Mapping System can be engineered so that a LISP mapping request can be serviced by a Map-Resolver that is close to the end-user. This approach reduces delays related to the processing of the “first packet”, which can be quite high with the legacy LISP control plane (at least equal to the RTT between the ITR and its MR). We propose two solutions to resolve this issue. The first solution consists in allowing the Mapping System to help forwarding packets that do not match an existing mapping record. The second solution is that the xTR prepares in advance the required mappings so that neither delay nor loss is experienced when receiving the first packet.

This framework advocates for a global Mapping System to be maintained locally. To that extent, we present hereafter new LISP primitives to allow for bulk retrieval of mappings and subscription to notifications when a predefined set of filters are hit.

IV. MAPPING SYSTEMS DISCOVERY AND ADVERTISEMENT

We present in the following sub-sections routing protocol extensions to dynamically advertise and discover Mapping Systems within and beyond a network domain.

A. A new LISP BGP community attribute

Because the design and operation of a consistent LISP Mapping System are critical for the adoption of the protocol at large scale, means to dynamically discover other Mapping Systems that are open to co-operate in inter-domain LISP deployment scenarios are required. A LISP domain may need to discover available Mapping Systems so that it can rely upon them to extend the reachability scope.

We propose to support the discovery of LISP Mapping Systems, deployed in distinct administrative domains, with a specific Border Gateway Protocol (BGP) community attribute [**bgpExtended**]. The detailed format of the new BGP community is described in [**lisplab**]. An advantage of adopting a BGP community attribute is that Mapping System interconnection functions can be integrated in standard BGP decision-process filters; on the other hand, a disadvantage is that a current practice is to filter out all the unknown BGP community attributes. Standardising this BGP Extended Communities will help this announcement to be safely propagated.

This BGP Extended Communities attribute is used to inform other domains about the support of the mapping service. EID that can be serviced with LISP will be tagged accordingly. Note that an EID can be serviced by multiple Mapping Systems. Remote LISP Mapping Systems will rely upon that BGP-based advertising capability to discover the existence and the status of other Mapping Systems.

Once a Mapping System is discovered, a local Mapping System can solicit the remote Mapping System to enter negotiation discussions for the establishment of an interconnection agreement with that remote Mapping System. The contact IP address provided as part of the BGP Extended Communities attribute will be used to contact a remote Mapping System to request for further LISP-related capabilities, possibly negotiate an interconnection agreement and, consequently, extend the scope of the networks that can be serviced using LISP. Also, leaf LISP-aware networks can rely upon the information carried in the BGP Extended Communities attribute to discover Mapping Systems that may be solicited to invoke their mapping service. Subscription cycles may then be considered.

B. A new interior gateway protocol feature

This section focuses on extensions to link-state routing protocols for the discovery and advertisement of LISP Mapping Service functions, especially the Map-Resolver and Map-Server LISP components within a domain. For example, such approach can use an extension of the Open Shortest Path First (OSPF) protocol. Such discovery allows for automatic operations of LISP networks. Mapping Service reachability information is announced into the domain by a router that embeds a Mapping Service Function instance, or which has been instructed (by means of specific configuration tasks, for example) to advertise such information on behalf of a third-party Mapping Service Function.




Fig. 4. Process to discover MS components: an example with OSPF.

The proposed mechanism may be used to advertise and learn MSFs that are available in the same administrative domain than xTRs. It can also be used to dynamically advertise related reachability information learned using other means when the MSFs and xTRs do not belong to the same administrative entity. To do so, a new Type-Length-Value (TLV)-encoded attribute, named the Mapping Service Function Discovery (MSFD) TLV, is defined. This attribute is carried in an OSPF Router Information Link State Advertisements (LSA). More details on the TLV attribute can be found in [lisplab].

The location of each Mapping Service Function is then flooded into the routing domain, as represented in Fig. 4 (considering the case the LSA is AS-scoped). The xTR routers deployed within the OSPF domain must listen to the flooding messages sent by active Mapping Service Function instances.

The information to be announced by means of the MSFD TLV carried in the LSA, during the LISP Mapping Service Function Discovery procedure, includes (but is not necessarily limited to): MSF type (Map-Resolver, Map-Server, or both), MSF locators (one or several IPs), unavailability timer, reboot timer, MSF diagnosis support, mapping database status, MSF status (Enabled, Disabled). All but the first two items are optional and may therefore be included in the Mapping Service Function Discovery messages. Additional capabilities such as the support of mapping bulk retrieval or notifications may also be advertised.

V. NEGOTIATION, INTERCONNECT AND INVOCATION

Let us present the control plane extensions to support the LISP-MSX negotiation, interconnection and invocation blocks, as illustrated by Fig. 5.

A. NEGOTIATION CYCLE

The proposal is to conduct the inter-Mapping System negotiation cycle by means of CPNP (Connectivity Provisioning Negotiation Protocol) [CPNP]. CPNP is meant to dynamically exchange

and negotiate the connectivity provisioning parameters between two LISP Mapping Systems. CPNP is used as a tool to introduce automation in the negotiation procedure, thereby fostering the overall mapping service delivery process. CPNP can be used to negotiate the parameters to connect two Mapping Systems or subscribe to services offered by a given Mapping System. With CPNP, each agreement can be identified by a local identifier (the CUSTOMER_AGREEMENT_IDENTIFIER) assigned by a local Mapping System but also with a unique identifier (the PROVIDER_AGREEMENT_IDENTIFIER) assigned by a peer Mapping System.

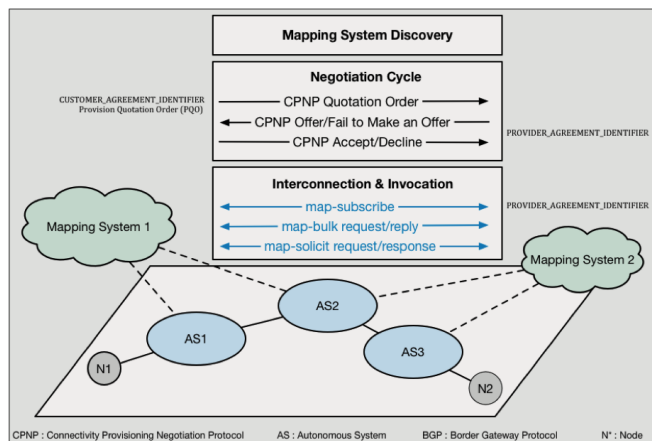


Fig. 5. An example of CPNP-based negotiation cycle and new LISP primitives used for the interconnection and invocation phases.

CPNP accommodates both technical and business-related requirements. Indeed, it supports various negotiation modes, including administrative validation operations. In particular, CPNP adopts a Quotation Order/Offer/Answer model where: (1) the Client specifies its requirements via a Provision Quotation Order (PQO), (2) the Server makes an offer to either address the requirements of the PQO or suggests a counter-proposal that partially addresses the requirements of the PQO or declines the PQO, then (3) the Client either accepts or declines the offer. Fig. 5 shows typical CPNP negotiation cycles. The PROVIDER_AGREEMENT_IDENTIFIER that is returned during the negotiation phase may be included in service invocation messages to ease correlating requests within a negotiation context (e.g., CPNP context; particularly, its integration in a Map-Request or a Map-Reply requires some modifications to the message formats.

B. NOVEL CONTROL PLANE PRIMITIVES

New LISP control plane primitives are defined to support the subscription and interconnection to Mapping Services, let alone their serviceability:

- Map-Subscribe/Map-Subscribe-Ack messages are exchanged between Mapping Services, possibly including a number of mapping filters

that the Mapping Service could support to trigger notifications to maintain the entries of the mapping database; the mapping “filter” is a novel feature of the proposed control plane primitives. A filter is used to transport any useful information, like flow and AS identifiers, for instance.

- Map-Bulk-Request/Map-Bulk-Reply messages are used to bypass the limitation of current LISP control plane primitives as far as the dissemination of mapping information is concerned. They allow to query multiple EID-prefixes with a single mapping request message by exploiting the mapping filter. In practice, the whole mapping database can be retrieved by exchanging one Map-Bulk-Request and as many Map-Bulk-Reply.
- Map-Solicit-Request messages are used, in the proposed framework, to enhance the robustness of LISP networks during such ITR failure events. While recovering from a failure, an ITR sends a Map-Solicit-Request to discovery other ITRs in the same routing domain. Upon receipt of the Map-Solicit-Request, another ITR replies with a Map-Solicit-Response message. With this process, the ITR has a list of peer ITRs, thanks to this Map-Bulk-Request/Reply signaling that runs between local xTRs to retrieve a copy of their mapping caches.

These features are detailed in [lisplab]. It is worth mentioning that these novel control-plane primitives are not primarily meant to replace existing basic LISP control plane primitives. Rather, they are meant to extend the LISP control plane behaviour in order to make LISP meeting the network management expectations of Internet Services and Network Providers more easily.

VI. EXPERIMENTAL RESULTS

We implemented the LISP-MSX solution [https://github.com/lip6-lisp] extending the LIP6-LISP OpenLISP control plane to support the new xTR and MS features and the Quagga router to include the new TLVs in both BGP and OSPF daemon [openlisp]. Then, we evaluated it within the LISP-LAB testbed [http://www.lisp-lab.org] and part of Cisco LISP Beta Network [http://www.lisp4.net] ; the former is an experimental platform, solely leveraging on LIP6-LISP OpenLISP nodes for all functions (xTR, MS/MR, PxTR), while the latter is using proprietary devices to run the control-plane and is managed by Cisco.

The LISP-lab mapping system is connected to the ‘LISP4.net’ mapping system via the DDT roots Lambda and Omega operated by LIP6, in Paris, France, and CSUC, in Barcelona, Spain, respectively. The roots have therefore a view on both mapping systems and are able to redirect resolution requests to authoritative MSes.

We use two LISP sites, one in LIP6 facility with one xTR, and another one in the LyonIX facility, Lyon, France, with another xTR. Note that each

LISP site belongs to a different Mapping System. One standard MS and one LISP-MSX MS are located in LyonIX. We deployed three MRs in the LIP6 site; while the first one (MR1) utilizes the LISP-MSX MS for handling mapping resolutions, the two others run the DDT protocol with the Lambda root located in LIP6 for one MR (MR2), and with the Omega root for the other MR (MR3). The resulting interconnection between the two Mapping Systems using LISP-MSX, the MR1 in the first Mapping System can discover the EID-prefix space of the second Mapping System and obtain related mapping entries. Hence, the MR1 can query directly the LISP-MSX MS, while MR2 and MR3 use the traditional DDT root system to resolve the EID-prefix belonging to the second Mapping System.




Fig. 6. Mapping resolution latency results over the LISP-LAB testbed (with logarithmic scales).

In our measurements, the LIP6 xTR acts as an ITR and the LyonIX xTR as an ETR. The ETR registers the same EID-prefix with both MSes. Therefore, Fig. 6 reports the time required to retrieve a mapping entry from the Mapping System in three scenarios:

- LISP-MSX: the proposed framework.
- Nearby DDT root, i.e., the Lambda root.
- Distant DDT root, i.e., the Omega root.

About seven hundred mapping resolutions (i.e., Map-Requests followed by Map-Replies) were executed for each case during three days.

For each measurement, the ITR in the LIP6 site sends the same Map-Request to the three MRs; we recorded the time when the Map-Request leaves the ITR, and the time when the Map-Reply message from the MS is received by the MR, hence computing the mapping resolution latency by subtracting the first from the second. Therefore, the difference in mapping resolution latency only depends on the time when the Map-Request leaves the MR, and the time when that Map-Request message is received by the MS. In the simplified LISP-MSX scenario, the Map-Request message is forwarded directly from MR in the LIP6 to the MS in LyonIX. While in

the last two cases, the MR uses DDT and so Map-Request messages are sent to the DDT roots and then reach the MS of the destination EID-prefix, in the LISP-MSX case DDT roots are bypassed and the Map-Request messages directly reach the MS.

The results in Fig. 6 shows that our framework can dramatically reduce the mapping resolution time, even compared to the mapping resolution service provided by the nearby DDT root, from a median around 5 ms with LISP-MSX to a median of about 25 ms with the nearby DDT root. This difference is explained by the forwarding stretch suffered by Map-Requests having to pass through the DDT root. Note that even in case (i) the DDT root is located in the same local network and (ii) the next DDT node is directly the destination MS, as in our setting, the MR still needs to query the DDT root. Instead, with LISP-MSX the Map-Request message is forwarded directly to the destination MS thanks to the pre-established interconnection. Therefore, in the worst case with a distant DDT root, the latency is increased even more, at least by a factor equal to the round-trip time between the source MR and the DDT root. An in-depth presentation is provided in a demo-tutorial video .

VII. PERSPECTIVES

LISP is a promising protocol to improve the forwarding of Internet traffic while mastering the growth of routing tables. Yet, it failed to be massively adopted so far, partly because of the operation of its Mapping System that may undesirably delay forwarding decisions at the cost of jeopardizing the performance of the LISP connectivity service.

This article discussed the LISP-MSX framework meant to improve LISP operation at the Internet scale, by facilitating cooperation between Mapping Systems and introducing more automation in the inter-domain connectivity service delivery procedure.

We believe such optimization could raise awareness among the service providers' community, yielding new business opportunities such as the monetization of LISP mapping services and the enforcement of advanced, service-inferred, inter-domain traffic engineering policies for the sake of better and strict QoS guarantees.

Mohamed Boucadair is a Senior IP Architect at Orange. Email: mohamed.boucadair@orange.com

Christian Jacquenet is the director of the strategic program office for advanced IP networking at Orange. Email: christian.jacquenet@orange.com

Dung Chi Phung is research engineer at Sorbonne Université, CNRS LIP6, France. Email: Chi-Dung.Phung@sorbonne-universite.fr

Stefano Secci is professor of networking at Cnam, Paris, France. Email: seccis@cnam.fr