LISP-MSX: Decentralized Interconnection of Independent LISP Mapping Systems

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

Abstract

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

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, that is, 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 years: there were approximately 10,000 IPv4 routes in 1994, and there are now more than 700,000 such routes. Likewise, there were a few hundred 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 of 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 growing interest within the Internet community. It is generally admitted that the ability to separate the "where" from the "who" allows getting 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 have been discussed in the last two decades, as documented in [1]. Among them, the Locator/ID Separation Protocol (LISP) is differentiated from most of the other approaches in that it does not imply any modification of terminals and underwent standardization for several years [2].

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 [3]. Several LISP mapping database systems were proposed, but the Delegated Database Tree (LISP-DDT) [4] is the one currently deployed by operational implementations (e.g., the Cisco IOS). LISP-DDT proposes a hierarchical resolution model like the Domain Name Service (DNS) 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 one resolver node to another one, passing through a DDT until it reaches an authoritative server. Alternative proposals were discussed, such as Alternative LISP Topology (ALT) [5], which, however, mandates a parallel node-disjoint separation for the control plane, with distinct Border Gateway Protocol (BGP) 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 interconnect independent mapping systems without following a resolution hierarchy. More precisely, we specify and experiment with novel LISP control-plane primitives and route discovery protocol extensions in support of decentralization of the mapping system interconnection and resolution process, decreasing the mapping resolution latency and providing better scaling of the overall mapping system. About the first aspect, the motivation to adopt LISP-MSX is to avoid falling into a similar situation as the DNS one, with three countries controlling the whole DNS root domain, and many operators wishing to be free of such dependence. LISP-MSX supports decentralized governance of mapping systems, allowing for customized mapping system implementations within provider boundaries. The framework is based on new mechanisms to dynamically select remote LISP mapping systems, negotiate and establish interconnect agreements with them, and optimize connectivity service operation.

LISP Background and Operations

LISP operation relies on the manipulation of two identifiers: the routing locator (RLOC), assigned to network topology attachment points, and the endpoint identifiers (EIDs), assigned to terminal devices independent of the network topology. LISP forwarding uses mapping functions that associate EIDs with RLOCs as well as an encapsula-

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

Mohamed Boucadair, Christian Jacquenet, and Dung Phung are with Orange, France; Stefano Secci is with with Cnam, France.

Digital Object Identifier: 10.1109/MCOM.2018.1701323

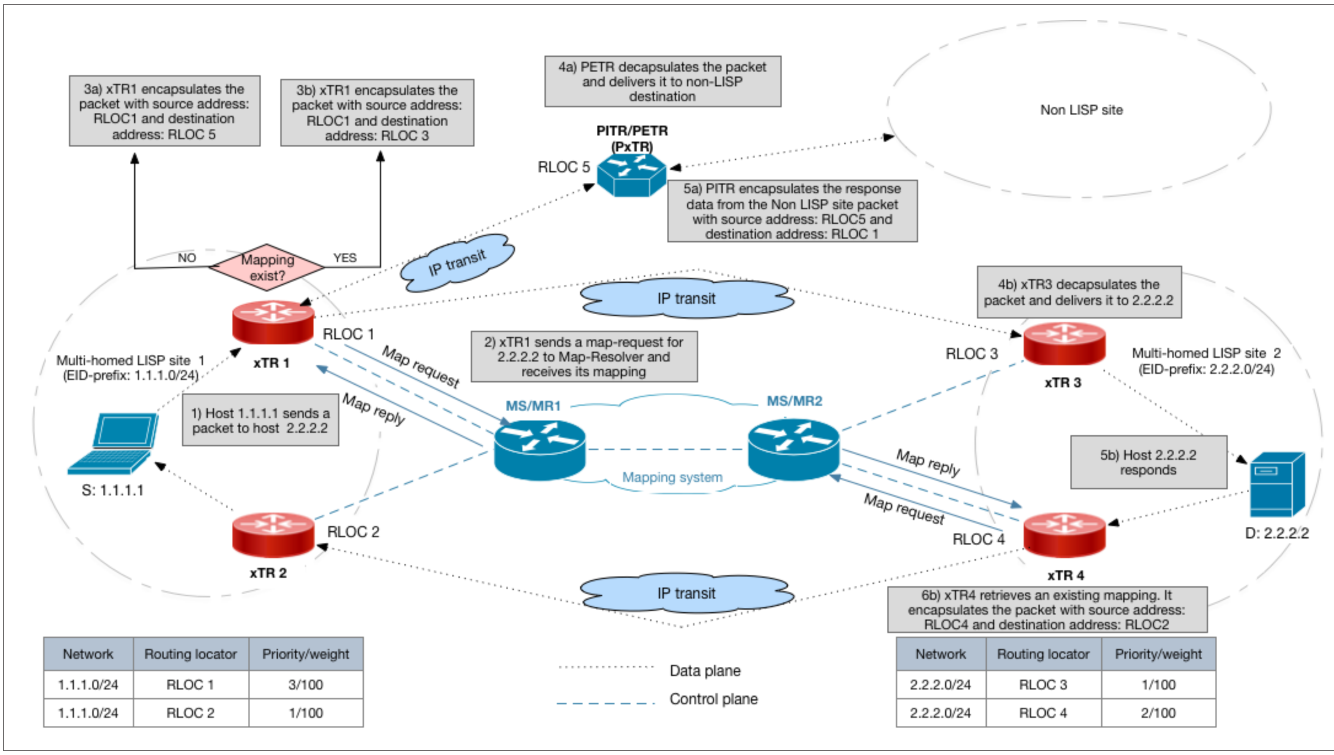

Figure 1. Example of LISP communications between two LISP sites.

tion scheme. As such, LISP does not mandate any specific modification of EID terminals: they use legacy IP addressing and forwarding. Routers that compose a LISP network are called tunnel routers. They are responsible for encapsulating/decapsulating 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 (MRs), Map servers (MSs), and other components like authorization and subscription servers are part of a mapping system. While an MS learns about authoritative EID-to-RLOC mappings from ETRs by means of Map-Register messages [2], and records them in the mapping database, an MR processes LISP Map-Requests [2] sent by ITRs and solicits MSs accordingly to resolve EID-to-RLOC mappings. The mapping resolution service of 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 (PxTRs) [6] are deployed to handle traffic between non-LISP and LISP sites. As such, a PITR behaves as an ITR on behalf of non-LISP sites that send packets to destinations located in LISP sites. Likewise, a PETR 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, host S in LISP site 1 (EID 1.1.1.1) has to communicate with 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 RLOCs of the destination by sending a Map-Request message to the Map-Resolver and getting a Map-Reply message. If the Map-Reply contains a positive mapping record, the packet is encapsulated by that ITR and forwarded towardsan 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 the case of many equal values, load balancing is done according to the weight metrics.) Then the ETR decapsulates the packet and forwards it natively to D. Note that if no EID-to-RLOC mapping is available in the mapping system, there are two possibilities. The first possibility is that the packet is encapsulated by that ITR and forwarded toward 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.

Challenges of LISP Operations 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 [7]. Indeed, many proposals have been investigated in the past few years, including mobile core networks [8], software-defined networks [9], and prefix de-aggregation control practices [10], leading to independent mapping systems that may benefit from interconnecting with each other.

Furthermore, multiple mapping systems will coexist for other reasons, including to avoid country-centric governance, allow for various technologies to implement the mapping service, take advantage of new business opportunities, encourage service innovation, and so on. The lack of clear policies for the operation and management of the LISP mapping systems encourages such practices.

Moreover, because the LISP mapping system may provide service differentiation opportunities, IP access and transit providers may consider operating a (local) mapping system. Mapping service providers may consider deploying innovative services to their customers, for example, the maintenance of local caches, or the update of ITR mapping entries that match some criteria requested by a LISP-enabled network. Mapping service 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, and so on.

As represented in Fig. 2, a LISP mapping system may handle one or multiple prefixes that belong to one or multiple autonomous systems (ASs). Distinct flavors of mapping systems may be deployed; each may rely on specific technology. As such, an interface to facilitate interconnection between these realms is to be specified.

A hierarchy in the mapping system organization for governance, regulatory, and business 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 on the availability of an inter-domain mapping system that spans several domains. From this perspective, the success of global LISP adoption and deployment mainly depends on how LISP-enabled domains (e.g., an AS or a simple local area network) will graft to existing mapping systems, which can guarantee global reachability. 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.

A Framework for Improving LISP Operations on a 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, which we refer to as LISP-MSX.

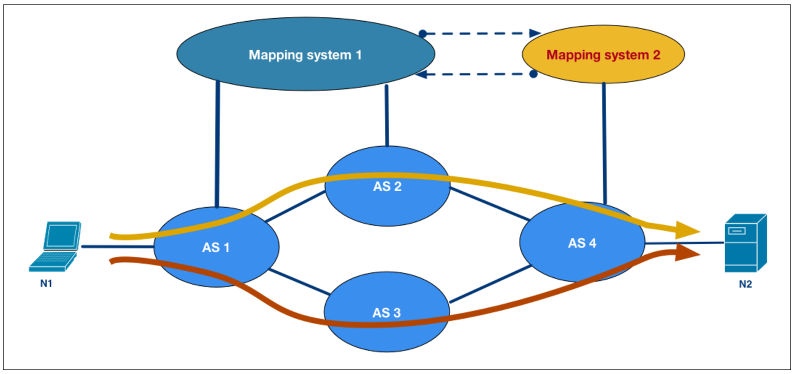

Figure 2. Example of LISP-MSX interconnection between two mapping systems.

An Interconnect Framework

In order to extend the reachability of LISP EIDs beyond the boundaries of a single mapping system, we aim to propose a framework that does not require changing xTR behavior such that an xTR would query multiple mapping systems concurrently (i.e., configured with multiple mapping servers of independent mapping systems). These mapping systems have to interconnect to extend their reachability and avoid pressure on PxTR devices. Moreover, mapping systems can 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 operation of a global mapping system at the scale of the Internet instead of a fragmented mapping system. Figure 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.

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 offered by one or more 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 for manual configurations as is the case for the current LISP specification, in particular for the configuration of MSs and DDT roots [11]. The goal of the Negotiation block is to negotiate inter-

The design and operation of a consistent LISP Mapping System are critical for the adoption of the protocol at large scale. Therefore, means to dynamically discover other Mapping Systems that are open to cooperate in inter-domain LISP deployment scenarios are required. To extend its reachability scope, a LISP domain may have to discover available Mapping Systems.

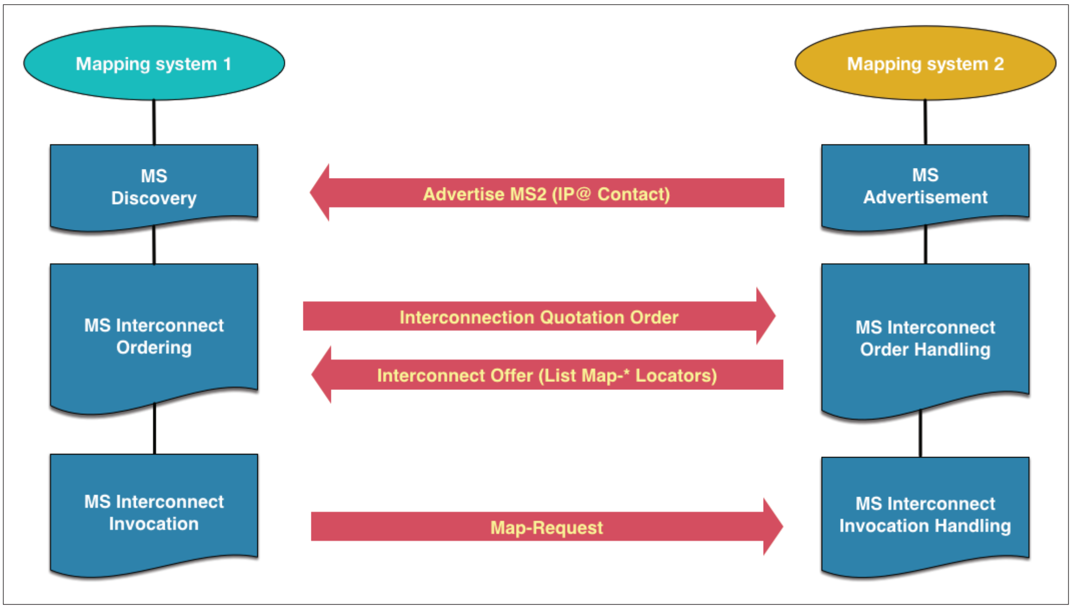

Figure 3. Representation of four functional blocks for LISP-MSX operations: Discovery, Negotiation, Interconnection, and Invocation.

connection 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. Each player involved in the operation of a LISP-based connectivity forwarding service needs to have a clear role in order to guide the definition of externals, and to help define and tune appropriate troubleshooting, diagnosis, and repair mechanisms. The inability to identify the root cause of a LISP connectivity degradation is a hurdle for deploying LISP on the scale of the Internet. The interconnection agreement can be unidirectional or bidirectional. Dedicated technical clauses may be included in the interconnect agreements to specify whether advanced primitives (e.g., bulk mapping transfer or record notifications) are supported, and also 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. All (or part) of the mapping table entries (or a part thereof) are exchanged between these mapping systems so that Map-Requests are processed close to the LISP leaf networks. Therefore, mapping resolution delays are shortened.

Service Invocation: Invoke a peer mapping system for mapping records resolution in particular. Other services can be offered by the mapping system, for example, assisting with the forwarding of the first packet before a mapping entry is available in the xTR cache.

Also, the mapping system can be engineered so that a LISP mapping request is serviced by a resolver close to the end user. First-packet processing delays are reduced with respect to the legacy LISP control plane (at least equal to the round-trip time, RTT, between the ITR and its MR). We propose two solutions for this issue. The first consists of allowing the mapping system to help forward packets that do not match an existing record; in the second, 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.

Mapping Systems Discovery and Advertisement

We present routing protocol extensions to dynamically advertise and discover mapping systems within and beyond a network domain.

A New LISP BGP Community Attribute

The design and operation of a consistent LISP mapping system are critical for the adoption of the protocol on a large scale. Therefore, means to dynamically discover other mapping systems that are open to cooperate in inter-domain LISP deployment scenarios are required. To extend its reachability scope, a LISP domain may have to discover available mapping systems.

We propose to support the discovery of LISP mapping systems, deployed in distinct administrative domains, with a specific BGP community attribute [12]. The detailed format of the new BGP community is described in [13]. 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 unknown BGP community attributes. Standardizing these 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. An 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 on that BGP-based advertising capability to discover the existence and 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 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 on 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.

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 an 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 may be announced into the domain by a router that embeds a Mapping Service Function (MSF) instance, or has been instructed by an operator.

The proposed mechanism may be used to advertise and learn MSFs that are available in the same administrative domain as xTRs. It can also be used to dynamically advertise associated 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 advertisement (LSA). More details on the TLV attribute can be found in [13]. The location of each MSF is disseminated within the domain as shown in Fig. 4.

The information disseminated using the MSFD TLV carried in the LSA includes: MSF type (Map-Resolver, Map-Server, or both), MSF locators (one or several IPs), unavailability timer, reboot timer, MSF diagnosis support, mapping database status, or MSF status (Enabled, Disabled). All but the first two items are optional and may therefore be included in the MSF Discovery messages. Other capabilities (e.g., the support of mapping bulk retrieval or notifications) may also be included in the MSFD TLV.

Negotiation, Interconnect, Invocation

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

Negotiation Cycle

The proposal is to conduct the inter-mapping system negotiation cycle by means of Connectivity Provisioning Negotiation Protocol (CPNP) [14].

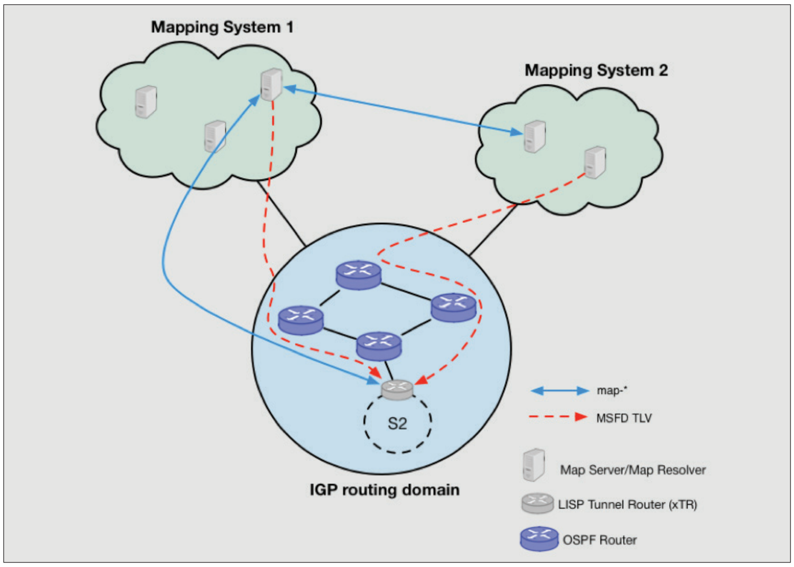

Figure 4. Process to discover MS components: an example with OSPF.

CPNP is meant to dynamically exchange and negotiate a set of connectivity parameters that are required to interconnect two 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.

CPNP accommodates both technical and business-related requirements. Indeed, it supports various negotiation modes, including administrative validation operations. In particular, CPNP adheres to the following model. The Client first asks for a quotation using a Provision Quotation Order (PQO) message that describes the expected service; the server can fully or partially satisfy the client's requirements; an offer is thus proposed to the Client. Alternatively, the Server may decline the quotation order. Last, the Client accepts or declines the offer. Figure 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; in particular, its integration in a Map-Request or a Map-Reply requires some modifications to the message formats.

Novel Control Plane Primitives

New LISP control plane primitives are defined to support the subscription and interconnection to mapping services. These primitives also allow increasing the serviceability of mapping services.

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

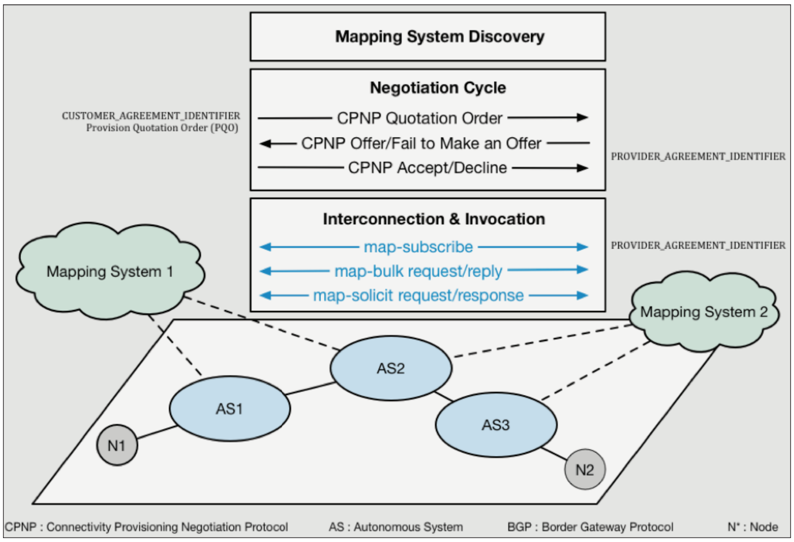

Figure 5. An example of the CPNP-based negotiation cycle and new LISP primitives used for the interconnection and invocation phases.

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 querying 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 messages as needed.

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 discover other ITRs in the same routing domain. Upon receipt of a Map-Solicit-Request, another ITR replies with a Map-Solicit-Response message. Through 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 [13]. 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 behavior in order to make LISP meet the network management expectations of Internet service and network providers more easily.

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 the BGP and OSPF daemons [15]. Then we evaluated it within the LISP-LAB testbed (http://www.lisp-lab.org) and part of the 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 uses 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 therefore have a view of both mapping systems and are able to redirect resolution requests to authoritative MSs.

We use two LISP sites, one in the 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 at the LIP6 site; while the first one (MR1) utilizes the LISP-MSX MS for handling mapping resolutions, the other two 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). With the resulting interconnection between the two mapping systems using LISP-MSX, MR1 in the first mapping system can discover the EID-prefix space of the second mapping system and obtain related mapping entries. Hence, MR1 can directly query 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.

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 MSs. Therefore, Fig. 6 reports the time required to retrieve a mapping entry from the mapping system in three scenarios:

1. LISP-MSX: the proposed framework
2. Nearby DDT root, that is, the Lambda root
3. Distant DDT root, that is, the Omega root

About 700 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 the MR in LIP6 to the MS in LyonIX. While in the last two cases the MR uses DDT, and hence 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 (i) even in the first case 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 RTT between the source MR and the DDT root. An in-depth presentation is provided in a demo-tutorial video.

Perspectives

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

We discuss the LISP-MSX framework, meant to improve LISP operation on 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 provider 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 quality of service guarantees.

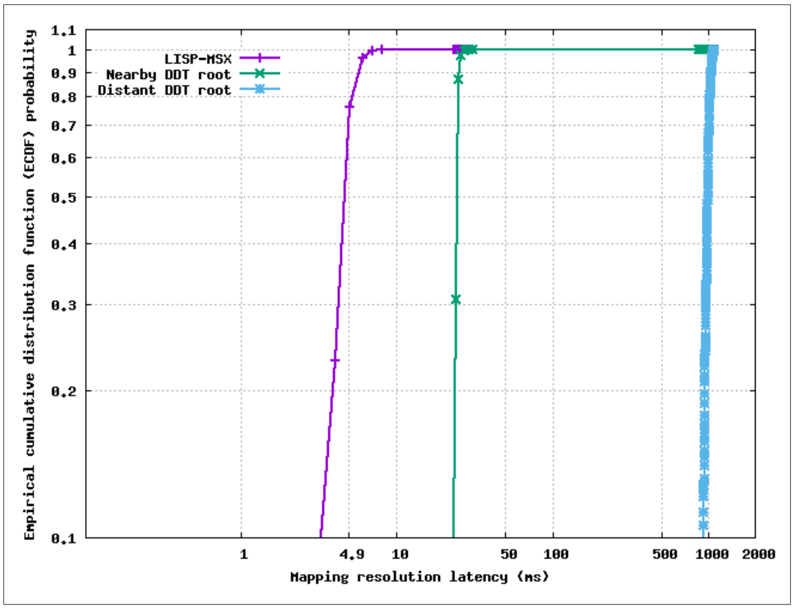

Figure 6. Mapping resolution latency results over the LISP-LAB testbed (with logarithmic scales).

References

[1] M. Komua, M. Sethia, and N. Beijara, "A Survey of Identifier-Locator Split Addressing Architectures," *Computer Science Review*, vol. 17, 2015, pp. 25–42.

[2] D. Farinacci *et al.*, "The Locator/Id Separation Protocol (LISP)," IETF RFC 6830, 2013.

[3] M. Hoefling, M. Menth, and M. Hartmann, "A Survey of Mapping Systems for Locator/Identifier Split Internet Routing," *IEEE Commun. Surveys & Tutorials*, vol. 15, 2013.

[4] V. Fuller *et al.*, "Locator/Id Separation Protocol Delegated Database Tree (LISP-DDT)," IETF RFC 8111, 2017.

[5] V. Fuller *et al.*, "Locator/Id Separation Protocol Alternative Logical Topology (LISP+ALT)," IETF RFC 6836, 2013.

[6] D. Lewis *et al.*, "Interworking between Locator/Id Separation Protocol (LISP) and Non-LISP Sites," IETF RFC 6832, 2013.

[7] T. Jeong *et al.*, "Experience on the Development of LISP-Enabled Services: An ISP Perspective," *Proc. IEEE NetSoft*, 2015, pp. 1–9.

[8] M. Portoles-Comeras *et al.*, "An Evolutionary Path for the Evolved Packet System," *IEEE Commun. Mag.*, vol. 53, 2015, pp. 184–91.

[9] A. Rodriguez-Nata *et al.*, "LISP: A Southbound SDN Protocol?" *IEEE Commun. Mag.*, vol. 53, 2015, pp. 201–07.

[10] M. Yannuzzi *et al.*, "Managing Interdomain Traffic in Latin America: A New Perspective Based on LISP," *IEEE Commun. Mag.*, vol. 47, 2009, pp. 40–48.

[11] R. Sambasivan *et al.*, "Bootstrapping Evolvability for Inter-Domain Routing," *Proc. ACM HotNets*, 2015.

[12] S. Sangli, D. Tappan, and Y. Rekhter, "BGP Extended Communities Attribute," IETF RFC 4360, 2006.

[13] "LISP-Lab Project Publications and IETF Documents," http://www.lisp-lab.org/publications.html, 2018, accessed Dec. 3, 2018.

[14] M. Boucadair *et al.*, "Connectivity Provisioning Negotiation Protocol (CPNP) draftboucadair-connectivity-provisioning-protocol-12, 2016.

[15] D. C. Phung *et al.*, "The OpenLISP Control-Plane Architecture," *IEEE Network*, vol. 38, 2004, pp. 34–40.

Biographies

Mohamed Boucadair (mohamed.boucadair@orange.com) is a senior IP architect at Orange, France, and an expert on VoIP services, IP/MPLS core networks, dynamic provisioning, and inter-domain traffic engineering.

Christian Jacquenet (christian.jacquenet@orange.com) joined Orange in 1989, and he is currently the Referent Expert of the "Networks of the Future" Orange Expert community. He is also the head of Orange's IPv6 strategic program. His research activities are in the areas of IP networking and network automation techniques, including software-defined networking.

Dung Chi Phung (Chi-Dung.Phung@orange.com) is a research engineer at Orange. He was a research engineer at LIP6, UPMC, from 2013 to 2018. He received his Ph.D. from UPMC in 2018. His research interests are Internet protocol design and experimentation.

Stefano Secci (stefano.secci@cnam.fr) became a professor of networking at Cnam, Paris, France, in 2018. Before that, he was an associate professor at LIP6, UPMC, from 2010 to 2018. He holds a Ph.D. from Telecom PariTech, France, and Politecnico di Milano, Italy. His research interests are about network virtualization and Internet routing.