

Internet Engineering Task Force (IETF)
Request for Comments: 7215
Category: Experimental
ISSN: 2070-1721

L. Jakab
Cisco Systems
A. Cabellos-Aparicio
F. Coras
J. Domingo-Pascual
Technical University of Catalonia
D. Lewis
Cisco Systems
April 2014

Locator/Identifier Separation Protocol (LISP)
Network Element Deployment Considerations

Abstract

This document is a snapshot of different Locator/Identifier Separation Protocol (LISP) deployment scenarios. It discusses the placement of new network elements introduced by the protocol, representing the thinking of the LISP working group as of Summer 2013. LISP deployment scenarios may have evolved since then. This memo represents one stable point in that evolution of understanding.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for examination, experimental implementation, and evaluation.

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

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7215>.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Tunnel Routers	5
2.1. Deployment Scenarios	5
2.1.1. Customer Edge (CE)	5
2.1.2. Provider Edge (PE)	6
2.1.3. Tunnel Routers behind NAT	8
2.1.3.1. ITR	8
2.1.3.2. ETR	9
2.1.3.3. Additional Notes	9
2.2. Functional Models with Tunnel Routers	9
2.2.1. Split ITR/ETR	9
2.2.2. Inter-Service-Provider Traffic Engineering	11
2.3. Summary and Feature Matrix	13
3. Map-Servers and Map-Resolvers	14
3.1. Map-Servers	14
3.2. Map-Resolvers	16
4. Proxy Tunnel Routers	17
4.1. PITRs	17
4.2. PETRs	18
5. Migration to LISP	19
5.1. LISP+BGP	19
5.2. Mapping Service Provider (MSP) PITR Service	20
5.3. Proxy-ITR Route Distribution (PITR-RD)	20
5.4. Migration Summary	23
6. Security Considerations	24
7. Acknowledgements	24
8. References	24
8.1. Normative References	24
8.2. Informative References	24
Appendix A. Step-by-Step Example BGP-to-LISP Migration Procedure ..	26
A.1. Customer Pre-Install and Pre-Turn-Up Checklist	26
A.2. Customer Activating LISP Service	28
A.3. Cut-Over Provider Preparation and Changes	29

1. Introduction

The Locator/Identifier Separation Protocol (LISP) is designed to address the scaling issues of the global Internet routing system identified in [RFC4984] by separating the current addressing scheme into Endpoint IDentifiers (EIDs) and Routing LOCators (RLOCs). The main protocol specification [RFC6830] describes how the separation is achieved and which new network elements are introduced, and it details the packet formats for the data and control planes.

LISP assumes that such separation is between the edge and core and uses mapping and encapsulation for forwarding. While the boundary between both is not strictly defined, one widely accepted definition places it at the border routers of stub autonomous systems, which may carry a partial or complete default-free zone (DFZ) routing table. The initial design of LISP took this location as a baseline for protocol development. However, the applications of LISP go beyond just decreasing the size of the DFZ routing table and include improved multihoming and ingress traffic engineering (TE) support for edge networks, and even individual hosts. Throughout this document, we will use the term "LISP site" to refer to these networks/hosts behind a LISP Tunnel Router. We formally define the following two terms:

Network element: Facility or equipment used in the provision of a communications service over the Internet [TELCO96].

LISP site: A single host or a set of network elements in an edge network under the administrative control of a single organization, delimited from other networks by LISP Tunnel Router(s).

Since LISP is a protocol that can be used for different purposes, it is important to identify possible deployment scenarios and the additional requirements they may impose on the protocol specification and other protocols. Additionally, this document is intended as a guide for the operational community for LISP deployments in their networks. It is expected to evolve as LISP deployment progresses, and the described scenarios are better understood or new scenarios are discovered.

Each subsection considers an element type and discusses the impact of deployment scenarios on the protocol specification. For definitions of terms, please refer to the appropriate documents (as cited in the respective sections).

This experimental document describes deployment considerations. These considerations and the LISP specifications have areas that require additional experience and measurement. LISP is not recommended for deployment beyond experimental situations. Results of experimentation may lead to modifications and enhancements of LISP mechanisms. Additionally, at the time of this writing there is no standardized security to implement. Beware that there are no countermeasures for any of the threats identified in [LISP-THREATS]. See Section 15 of [RFC6830] for specific known issues that are in need of further work during development, implementation, and experimentation, and see [LISP-THREATS] for recommendations to ameliorate the above-mentioned security threats.

2. Tunnel Routers

The device that is the gateway between the edge and the core is called a Tunnel Router (xTR); it performs one or both of two separate functions:

1. Encapsulating packets originating from an end host to be transported over intermediary (transit) networks towards the other endpoint of the communication.
2. Decapsulating packets entering from intermediary (transit) networks, originated at a remote end host.

The first function is performed by an Ingress Tunnel Router (ITR) and the second by an Egress Tunnel Router (ETR).

Section 8 of the main LISP specification [RFC6830] has a short discussion of where Tunnel Routers can be deployed and some of the associated advantages and disadvantages. This section adds more detail to the scenarios presented there and provides additional scenarios as well. Furthermore, this section discusses functional models, that is, network functions that can be achieved by deploying Tunnel Routers in specific ways.

2.1. Deployment Scenarios

2.1.1. Customer Edge (CE)

The first scenario we discuss is the customer edge, when xTR functionality is placed on the router(s) that connects the LISP site to its upstream(s) but is under its control. As such, this is the most common expected scenario for xTRs, and this document considers it the reference location, comparing the other scenarios to this one.




Figure 1: xTRs at the Customer Edge

From the LISP site's perspective, the main advantage of this type of deployment (compared to the one described in the next section) is having direct control over its ingress traffic engineering. This makes it easy to set up and maintain active/active, active/backup, or more complex TE policies, adding ISPs and additional xTRs at will, without involving third parties.

Being under the same administrative control, reachability information of all ETRs is easier to synchronize, because the necessary control traffic can be allowed between the locators of the ETRs. A correct synchronous global view of the reachability status is thus available, and the Locator-Status-Bits can be set correctly in the LISP data header of outgoing packets.

By placing the Tunnel Router at the edge of the site, existing internal network configuration does not need to be modified. Firewall rules, router configurations, and address assignments inside the LISP site remain unchanged. This helps with incremental deployment and allows a quick upgrade path to LISP. For larger sites distributed in geographically diverse points of presence (PoPs) and having many external connections and complex internal topology, it may, however, make more sense to both encapsulate and decapsulate as soon as possible, to benefit from the information in the IGP to choose the best path. See Section 2.2.1 for a discussion of this scenario.

Another thing to consider when placing Tunnel Routers is MTU issues. Encapsulation increases the amount of overhead associated with each packet. This added overhead decreases the effective end-to-end path MTU (unless fragmentation and reassembly are used). Some transit networks are known to provide larger MTU values than the typical value of 1500 bytes for popular access technologies used at end hosts (e.g., IEEE 802.3 and 802.11). However, placing the LISP router connecting to such a network at the customer edge could possibly bring up MTU issues, depending on the link type to the provider as opposed to the following scenario. See [RFC4459] for MTU considerations of tunneling protocols and how to mitigate potential issues. Still, even with these mitigations, path MTU issues are still possible.

2.1.2. Provider Edge (PE)

The other location at the core-edge boundary for deploying LISP routers is at the Internet service provider edge. The main incentive for this case is that the customer does not have to upgrade the CE router(s) or change the configuration of any equipment. Encapsulation/decapsulation happens in the provider's network, which may be able to serve several customers with a single device. For

large ISPs with many residential/business customers asking for LISP, this can lead to important savings, since there is no need to upgrade the software (or hardware, if that's the case) at each client's location. Instead, they can upgrade the software (or hardware) on a few PE routers serving the customers. This scenario is depicted in Figure 2.




Figure 2: xTRs at the Provider Edge

While this approach can make transition easy for customers and may be cheaper for providers, the LISP site loses one of the main benefits of LISP: ingress traffic engineering. Since the provider controls the ETRs, additional complexity would be needed to allow customers to modify their mapping entries.

The problem is aggravated when the LISP site is multihomed. Consider the scenario in Figure 2: whenever a change to TE policies is required, the customer contacts both ISP1 and ISP2 to make the necessary changes on the routers (if they provide this possibility). It is, however, unlikely that both ISPs will apply changes simultaneously, which may lead to inconsistent state for the mappings of the LISP site. Since the different upstream ISPs are usually competing business entities, the ETRs may even be configured to compete, to either attract all the traffic or get no traffic. The former will happen if the customer pays per volume, the latter if the connectivity has a fixed price. A solution could be to configure the Map-Server(s) to do proxy-replying and have the Mapping Service Provider (MSP) apply policies.

Additionally, since xTR1, xTR2, and xTR3 are in different administrative domains, locator reachability information is unlikely to be exchanged among them, making it difficult to set the Locator-Status-Bits (LSBs) correctly on encapsulated packets. Because of this, and due to the security concerns about LSBs as described in [LISP-THREATS], their use is discouraged (set the L-bit to 0). Map-Versioning is another alternative [RFC6834].

Compared to the customer edge scenario, deploying LISP at the provider edge might have the advantage of diminishing potential MTU issues, because the Tunnel Router is closer to the core, where links typically have higher MTUs than edge network links.

2.1.3. Tunnel Routers behind NAT

"NAT" in this section refers to IPv4 network address and port translation.

2.1.3.1. ITR




Figure 3: ITR behind NAT

Packets encapsulated by an ITR are just UDP packets from a NAT device's point of view, and they are handled like any UDP packet; there are no additional requirements for LISP data packets.

Map-Requests sent by an ITR, which create the state in the NAT table, have a different 5-tuple in the IP header than the Map-Reply generated by the authoritative ETR. Since the source address of this packet is different from the destination address of the request packet, no state will be matched in the NAT table and the packet will be dropped. To avoid this, the NAT device has to do the following:

- o Send all UDP packets with source port 4342, regardless of the destination port, to the RLOC of the ITR. The simplest way to achieve this is configuring 1:1 NAT mode from the external RLOC of the NAT device to the ITR's RLOC (called "DMZ" mode in consumer broadband routers).
- o Rewrite the ITR-AFI and "Originating ITR RLOC Address" fields in the payload.

This setup supports only a single ITR behind the NAT device.

2.1.3.2. ETR

An ETR placed behind NAT is reachable from the outside by the Internet-facing locator of the NAT device. It needs to know this locator (and configure a loopback interface with it), so that it can use it in Map-Reply and Map-Register messages. Thus, support for dynamic locators for the mapping database is needed in LISP equipment.

Again, only one ETR behind the NAT device is supported.




Figure 4: ETR behind NAT

2.1.3.3. Additional Notes

An implication of the issues described above is that LISP sites with xTRs cannot be behind carrier-based NATs, since two different sites would collide on the same forwarded UDP port. An alternative to static hole-punching to explore is the use of the Port Control Protocol (PCP) [RFC6887].

We only include this scenario due to completeness, to show that a LISP site can be deployed behind NAT should it become necessary. However, LISP deployments behind NAT should be avoided, if possible.

2.2. Functional Models with Tunnel Routers

This section describes how certain LISP deployments can provide network functions.

2.2.1. Split ITR/ETR

In a simple LISP deployment, xTRs are located at the border of the LISP site (see Section 2.1.1). In this scenario, packets are routed inside the domain according to the EID. However, more complex networks may want to route packets according to the destination RLOC. This would enable them to choose the best egress point.

The LISP specification separates the ITR and ETR functionality and allows both entities to be deployed in separated network equipment. ITRs can be deployed closer to the host (i.e., access routers). This way, packets are encapsulated as soon as possible, and egress point selection is driven by operational policy. In turn, ETRs can be deployed at the border routers of the network, and packets are decapsulated as soon as possible. Once decapsulated, packets are routed based on the destination EID according to internal routing policy.

We can see an example in Figure 5. The Source (S) transmits packets using its EID, and in this particular case packets are encapsulated at ITR_1. The encapsulated packets are routed inside the domain according to the destination RLOC and can egress the network through the best point (i.e., closer to the RLOC's Autonomous System (AS)). On the other hand, inbound packets are received by ETR_1, which decapsulates them. Then, packets are routed towards S according to the EID, again following the best path.

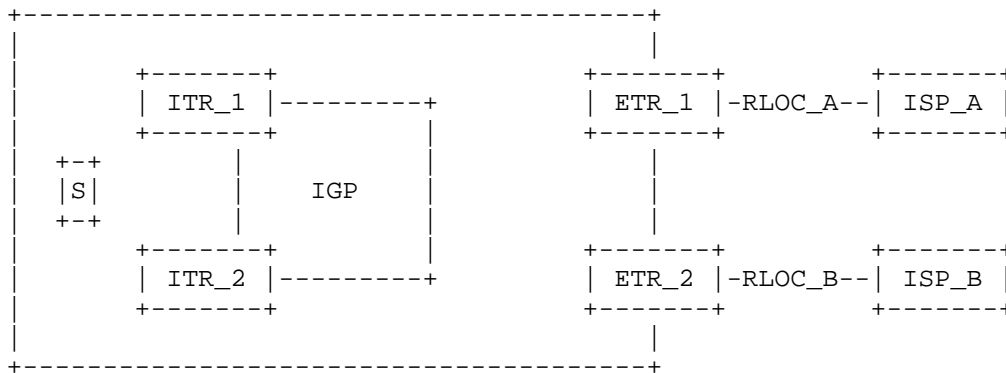


Figure 5: Split ITR/ETR Scenario

This scenario has a set of implications:

- o The site must carry more-specific routes in order to choose the best egress point, and typically BGP is used for this, increasing the complexity of the network. However, this is usually already the case for LISP sites that would benefit from this scenario.
- o If the site is multihomed to different ISPs and any of the upstream ISPs are doing unicast reverse path forwarding (uRPF) filtering, this scenario may become impractical. To set the correct source RLOC in the encapsulation header, ITRs need to first determine which ETR will be used by the outgoing packet. This adds complexity and reliability concerns.

- o In LISP, ITRs set the reachability bits when encapsulating data packets. Hence, ITRs need a mechanism to be aware of the liveness of all ETRs serving their site.
- o The MTU within the site network must be large enough to accommodate encapsulated packets.
- o In this scenario, each ITR is serving fewer hosts than in the case when it is deployed at the border of the network. It has been shown that the cache hit rate grows logarithmically with the amount of users [CACHE]. Taking this into account, when ITRs are deployed closer to the host the effectiveness of the mapping cache may be lower (i.e., the miss rate is higher). Another consequence of this is that the site may transmit a higher amount of Map-Requests, increasing the load on the distributed mapping database.
- o By placing the ITRs inside the site, they will still need global RLOCs. This may add complexity to intra-site routing configurations and more intra-site issues when there is a change of providers.

2.2.2. Inter-Service-Provider Traffic Engineering

At the time of this writing, if two ISPs want to control their ingress TE policies for transit traffic between them, they need to rely on existing BGP mechanisms. This typically means deaggregating prefixes to choose on which upstream link packets should enter. This either is not feasible (if fine-grained per-customer control is required, the very-specific prefixes may not be propagated) or increases DFZ table size.

Typically, LISP is seen as applicable only to stub networks; however, LISP can also be applied in a recursive manner, providing service provider ingress/egress TE capabilities without impacting the DFZ table size.

In order to implement this functionality with LISP, consider the scenario depicted in Figure 6. The two ISPs willing to achieve ingress/egress TE are labeled as ISP_A and ISP_B. They are servicing Stub1 and Stub2, respectively. Both are required to be LISP sites with their own xTRs. In this scenario, we assume that Stub1 and Stub2 are communicating with each other; thus, ISP_A and ISP_B offer transit for such communications. ISP_A has RLOC_A1 and RLOC_A2 as upstream IP addresses, while ISP_B has RLOC_B1 and RLOC_B2. The shared goal among ISP_A and ISP_B is to control the transit traffic flow between RLOC_A1/A2 and RLOC_B1/B2.




Figure 6: Inter-Service-Provider TE Scenario

Both ISPs deploy xTRs on RLOC_A1/A2 and RLOC_B1/B2, respectively and reach a bilateral agreement to deploy their own private mapping system. This mapping system contains bindings between the RLOCs of Stub1 and Stub2 (owned by ISP_A and ISP_B, respectively) and RLOC_A1/A2 and RLOC_B1/B2. Such bindings are in fact the TE policies between both ISPs, and the convergence time is expected to be fast, since ISPs only have to update/query a mapping to/from the database.

The packet flow is as follows. First, a packet originated at Stub1 towards Stub2 is LISP encapsulated by Stub1's xTR. The xTR of ISP_A recursively encapsulates it, and according to the TE policies stored in the private mapping system the ISP_A xTR chooses RLOC_B1 or RLOC_B2 as the outer encapsulation destination. Note that the packet transits between ISP_A and ISP_B double-encapsulated. Upon reception at the xTR of ISP_B, the packet is decapsulated and sent towards Stub2, which performs the last decapsulation.

This deployment scenario, which uses recursive LISP, includes three important caveats. First, it is intended to be deployed between only two ISPs. If more than two ISPs use this approach, then either the xTRs deployed at the participating ISPs must query multiple mapping systems, or the ISPs must agree on a common shared mapping system. Furthermore, keeping this deployment scenario restricted to only two ISPs maintains a scalable solution, given that only two entities need to agree on using recursive LISP and only one private mapping system is involved.

Second, the scenario is only recommended for ISPs providing connectivity to LISP sites, such that source RLOCs of packets to be recursively encapsulated belong to said ISP. Otherwise, the participating ISPs must register prefixes they do not own in the above-mentioned private mapping system. This results in either requiring complex authentication mechanisms or enabling simple traffic redirection attacks. Failure to follow these recommendations may lead to operational security issues when deploying this scenario.

And third, recursive encapsulation models are typically complex to troubleshoot and debug.

Besides these recommendations, the main disadvantages of this deployment case are:

- o An extra LISP header is needed. This increases the packet size and requires that the MTU between both ISPs accommodate double-encapsulated packets.
- o The ISP ITR must encapsulate packets and therefore must know the RLOC-to-RLOC bindings. These bindings are stored in a mapping database and may be cached in the ITR's mapping cache. Cache misses lead to an additional lookup latency, unless a push-based mapping system is used for the private mapping system.
- o Maintaining the shared mapping database involves operational overhead.

2.3. Summary and Feature Matrix

When looking at the deployment scenarios and functional models above, there are several things to consider when choosing an appropriate model, depending on the type of the organization doing the deployment.

For home users and small sites that wish to multihome and have control over their ISP options, the "CE" scenario offers the most advantages: it's simple to deploy, and in some cases it only requires a software upgrade of the Customer Premises Equipment (CPE), getting mapping service, and configuring the router. It retains control of TE and choosing upstreams by the user. It doesn't provide too many advantages to ISPs, due to the lessened dependence on their services in cases of multihomed clients. It is also unlikely that ISPs wishing to offer LISP to their customers will choose the "CE" model, as they would need to send a technician to each customer and, potentially, a new CPE device. Even if they have remote control over the router and a software upgrade could add LISP support, the operation is too risky.

For a network operator, a good option to deploy is the "PE" scenario, unless a hardware upgrade is required for its edge routers to support LISP (in which case upgrading CPEs may be simpler). It retains control of TE as well as the choice of Proxy Egress Tunnel Router (PETR) and Map-Server/Map-Resolver. It also lowers potential MTU issues, as discussed above. Network operators should also explore the "inter-service-provider TE" (recursive) functional model for their TE needs.

To optimize their traffic flow, large organizations can benefit the most from the "split ITR/ETR" functional model.

The following table gives a quick overview of the features supported by each of the deployment scenarios discussed above (marked with an "x" in the appropriate column): "CE" for customer edge, "PE" for provider edge, "Split" for split ITR/ETR, and "Recursive" for inter-service-provider traffic engineering. The discussed features include:

Control of ingress TE: This scenario allows the LISP site to easily control LISP ingress traffic engineering policies.

No modifications to existing int. network infrastructure: This scenario doesn't require the LISP site to modify internal network configurations.

Locator-Status-Bits sync: This scenario allows easy synchronization of the Locator Status Bits.

MTU/PMTUD issues minimized: The scenario minimizes potential MTU and Path MTU Discovery (PMTUD) issues.

Feature	CE	PE	Split	Recursive	NAT
Control of ingress TE	x	-	x	x	x
No modifications to existing int. network infrastructure	x	x	-	-	x
Locator-Status-Bits sync	x	-	x	x	-
MTU/PMTUD issues minimized	-	x	-	-	-

3. Map-Servers and Map-Resolvers

Map-Servers and Map-Resolvers make up the LISP mapping system and provide a means to find authoritative EID-to-RLOC mapping information, conforming to [RFC6833]. They are meant to be deployed in RLOC space, and their operation behind NAT is not supported.

3.1. Map-Servers

The Map-Server learns EID-to-RLOC mapping entries from an authoritative source and publishes them in the distributed mapping database. These entries are learned through authenticated Map-Register messages sent by authoritative ETRs. Also, upon reception of a Map-Request, the Map-Server verifies that the destination EID matches an EID-Prefix for which it is authoritative and then re-encapsulates and forwards it to a matching ETR. Map-Server functionality is described in detail in [RFC6833].

The Map-Server is provided by a Mapping Service Provider (MSP). The MSP participates in the global distributed mapping database infrastructure by setting up connections to other participants according to the specific mapping system that is employed (e.g., Alternative Logical Topology (ALT) [RFC6836], Delegated Database Tree (DDT) [LISP-DDT]). Participation in the mapping database and the storing of EID-to-RLOC mapping data are subject to the policies of the "root" operators, who should check ownership rights for the EID-Prefixes stored in the database by participants. These policies are out of scope for this document.

The LISP DDT protocol is used by LISP MSPs to provide reachability between those providers' Map-Resolvers and Map-Servers. The DDT root is currently operated by a collection of organizations on an open basis. See [DDT-ROOT] for more details. Similarly to the DNS root, it has several different server instances using names of the letters of the Greek alphabet (alpha, delta, etc.), operated by independent organizations. When this document was published, there were 6 such instances, with one of them being anycasted. [DDT-ROOT] provides the list of server instances on its web site and configuration files for several Map-Server implementations. The DDT root and LISP Mapping Providers both rely on and abide by existing allocation policies as defined by Regional Internet Registries (RIRs) to determine prefix ownership for use as EIDs.

It is expected that the DDT root organizations will continue to evolve in response to experimentation with LISP deployments for Internet edge multihoming and VPN use cases.

In all cases, the MSP configures its Map-Server(s) to publish the prefixes of its clients in the distributed mapping database and start encapsulating and forwarding Map-Requests to the ETRs of the AS. These ETRs register their prefix(es) with the Map-Server(s) through periodic authenticated Map-Register messages. In this context, for some LISP sites, there is a need for mechanisms to:

- o Automatically distribute EID-Prefix(es) shared keys between the ETRs and the EID-registrar Map-Server.
- o Dynamically obtain the address of the Map-Server in the ETR of the AS.

The Map-Server plays a key role in the reachability of the EID-Prefixes it is serving. On one hand, it is publishing these prefixes into the distributed mapping database, and on the other hand, it is encapsulating and forwarding Map-Requests to the authoritative ETRs of these prefixes. ITRs encapsulating towards EIDs for which a failed Map-Server is responsible will be unable to

look up any of their covering prefixes. The only exceptions are the ITRs that already contain the mappings in their local caches. In this case, ITRs can reach ETRs until the entry expires (typically 24 hours). For this reason, redundant Map-Server deployments are desirable. A set of Map-Servers providing high-availability service to the same set of prefixes is called a redundancy group. ETRs are configured to send Map-Register messages to all Map-Servers in the redundancy group. The configuration for fail-over (or load-balancing, if desired) among the members of the group depends on the technology behind the mapping system being deployed. Since ALT is based on BGP and DDT takes its inspiration from the Domain Name System (DNS), deployments can leverage current industry best practices for redundancy in BGP and DNS. These best practices are out of scope for this document.

Additionally, if a Map-Server has no reachability for any ETR serving a given EID block, it should not originate that block into the mapping system.

3.2. Map-Resolvers

A Map-Resolver is a network infrastructure component that accepts LISP-encapsulated Map-Requests, typically from an ITR, and finds the appropriate EID-to-RLOC mapping by consulting the distributed mapping database. Map-Resolver functionality is described in detail in [RFC6833].

Anyone with access to the distributed mapping database can set up a Map-Resolver and provide EID-to-RLOC mapping lookup service. Database access setup is mapping system specific.

For performance reasons, it is recommended that LISP sites use Map-Resolvers that are topologically close to their ITRs. ISPs supporting LISP will provide this service to their customers, possibly restricting access to their user base. LISP sites not in this position can use open access Map-Resolvers, if available. However, regardless of the availability of open access resolvers, the MSP providing the Map-Server(s) for a LISP site should also make available Map-Resolver(s) for the use of that site.

In medium- to large-size ASes, ITRs must be configured with the RLOC of a Map-Resolver; this type of operation can be done manually. However, in Small Office/Home Office (SOHO) scenarios, a mechanism for autoconfiguration should be provided.

One solution to avoid manual configuration in LISP sites of any size is the use of anycast [RFC4786] RLOCs for Map-Resolvers, similar to the DNS root server infrastructure. Since LISP uses UDP

encapsulation, the use of anycast would not affect reliability. LISP routers are then shipped with a preconfigured list of well-known Map-Resolver RLOCs, which can be edited by the network administrator, if needed.

The use of anycast also helps improve mapping lookup performance. Large MSPs can increase the number and geographical diversity of their Map-Resolver infrastructure, using a single anycasted RLOC. Once LISP deployment is advanced enough, very large content providers may also be interested in running this kind of setup, to ensure minimal connection setup latency for those connecting to their network from LISP sites.

While Map-Servers and Map-Resolvers implement different functionalities within the LISP mapping system, they can coexist on the same device. For example, MSPs offering both services can deploy a single Map-Resolver/Map-Server in each PoP where they have a presence.

4. Proxy Tunnel Routers

4.1. PITRs

Proxy Ingress Tunnel Routers (PITRs) are part of the non-LISP/LISP transition mechanism, allowing non-LISP sites to reach LISP sites. They announce via BGP certain EID-Prefixes (aggregated, whenever possible) to attract traffic from non-LISP sites towards EIDs in the covered range. They do the mapping system lookup and encapsulate received packets towards the appropriate ETR. Note that for the reverse path, LISP sites can reach non-LISP sites by simply not encapsulating traffic. See [RFC6832] for a detailed description of PITR functionality.

The success of new protocols depends greatly on their ability to maintain backwards compatibility and interoperate with the protocol(s) they intend to enhance or replace, and on the incentives to deploy the necessary new software or equipment. A LISP site needs an interworking mechanism to be reachable from non-LISP sites. A PITR can fulfill this role, enabling early adopters to see the benefits of LISP, similar to tunnel brokers helping the transition from IPv4 to IPv6. A site benefits from new LISP functionality (proportionally with existing global LISP deployment) when migrating to LISP, so it has the incentives to deploy the necessary Tunnel Routers. In order to be reachable from non-LISP sites, it has two options: keep announcing its prefix(es) with BGP, or have a PITR announce prefix(es) covering them.

If the goal of reducing the DFZ routing table size is to be reached, the second option is preferred. Moreover, the second option allows LISP-based ingress traffic engineering from all sites. However, the placement of PITRs significantly influences performance and deployment incentives. Section 5 is dedicated to the migration to a LISP-enabled Internet and includes deployment scenarios for PITRs.

4.2. PETRs

In contrast to PITRs, PETRs are not required for the correct functioning of all LISP sites. There are two cases where they can be of great help:

- o LISP sites with unicast reverse path forwarding (uRPF) restrictions, and
- o Communication between sites using different address family RLOCs.

In the first case, uRPF filtering is applied at the LISP site's upstream provider's PE router. When forwarding traffic to non-LISP sites, an ITR does not encapsulate packets, leaving the original IP headers intact. As a result, packets will have EIDs in their source address. Since we are discussing the transition period, we can assume that a prefix covering the EIDs belonging to the LISP site is advertised to the global routing tables by a Pitr, and the PE router has a route towards it. However, the next hop will not be on the interface towards the CE router, so non-encapsulated packets will fail uRPF checks.

To avoid this filtering, the affected ITR encapsulates packets towards the locator of the PETR for non-LISP destinations. Now the source address of the packets, as seen by the PE router, is the ITR's locator, which will not fail the uRPF check. The PETR then decapsulates and forwards the packets.

The second use case is IPv4-to-IPv6 transition. Service providers using older access network hardware that only supports IPv4 can still offer IPv6 to their clients by providing a CPE device running LISP, and PETR(s) for accessing IPv6-only non-LISP sites and LISP sites, with IPv6-only locators. Packets originating from the client LISP site for these destinations would be encapsulated towards the PETR's IPv4 locator. The PETR is in a native IPv6 network, decapsulating and forwarding packets. For non-LISP destinations, the packet travels natively from the PETR. For LISP destinations with IPv6-only locators, the packet will go through a Pitr in order to reach its destination.

For more details on PETRs, see [RFC6832].

PETRs can be deployed by ISPs wishing to offer value-added services to their customers. As is the case with PITRs, PETRs too may introduce path stretch (the ratio between the cost of the selected path and that of the optimal path). Because of this, the ISP needs to consider the tradeoff of using several devices close to the customers to minimize it, or fewer devices farther away from the customers to minimize cost instead.

Since the deployment incentives for PITRs and PETRs are different, it is likely that they will be deployed in separate devices, except for the Content Delivery Network (CDN) case, which may deploy both in a single device.

In all cases, the existence of a PETER involves another step in the configuration of a LISP router. CPE routers, which are typically configured by DHCP, stand to benefit most from PETRs. Autoconfiguration of the PETER locator could be achieved by a DHCP option or by adding a PETER field to either Map-Notify or Map-Reply messages.

5. Migration to LISP

This section discusses a deployment architecture to support the migration to a LISP-enabled Internet. The loosely defined terms "early transition phase", "late transition phase", and "LISP Internet phase" refer to time periods when LISP sites are a minority, a majority, or represent all edge networks, respectively.

5.1. LISP+BGP

For sites wishing to migrate to LISP with their Provider-Independent (PI) prefix, the least disruptive way is to upgrade their border routers to support LISP and register the prefix into the LISP mapping system, but to keep announcing it with BGP as well. This way, LISP sites will reach them over LISP, while legacy sites will be unaffected by the change. The main disadvantage of this approach is that no decrease in the DFZ routing table size is achieved. Still, just increasing the number of LISP sites is an important gain, as an increasing LISP/non-LISP site ratio may decrease the need for BGP-based traffic engineering that leads to prefix deaggregation. That, in turn, may lead to a decrease in the DFZ size and churn in the late transition phase.

This scenario is not limited to sites that already have their prefixes announced with BGP. Newly allocated EID blocks could follow this strategy as well during the early LISP deployment phase, depending on the cost/benefit analysis of the individual networks. Since this leads to an increase in the DFZ size, the following architecture should be preferred for new allocations.

5.2. Mapping Service Provider (MSP) PITR Service

In addition to publishing their clients' registered prefixes in the mapping system, MSPs with enough transit capacity can offer PITR service to them as a separate service. This service is especially useful for new PI allocations to sites without existing BGP infrastructure wishing to avoid BGP altogether. The MSP announces the prefix into the DFZ, and the client benefits from ingress traffic engineering without prefix deaggregation. The downside of this scenario is added path stretch.

Routing all non-LISP ingress traffic through a third party that is not one of its ISPs is only feasible for sites with modest amounts of traffic (like those using the IPv6 tunnel broker services today), especially in the first stage of the transition to LISP, with a significant number of legacy sites. This is because the handling of said traffic is likely to result in additional costs, which would be passed down to the client. When the LISP/non-LISP site ratio becomes high enough, this approach can prove increasingly attractive.

Compared to LISP+BGP, this approach avoids DFZ bloat caused by prefix deaggregation for traffic engineering purposes, resulting in slower routing table increase in the case of new allocations and potential decrease for existing ones. Moreover, MSPs serving different clients with adjacent aggregatable prefixes may lead to additional decrease, but quantifying this decrease is subject to future research study.

5.3. Proxy-ITR Route Distribution (PITR-RD)

Instead of a LISP site or the MSP announcing its EIDs with BGP to the DFZ, this function can be outsourced to a third party, a PITR Service Provider (PSP). This will result in a decrease in operational complexity at both the site and the MSP.

The PSP manages a set of distributed PITR(s) that will advertise the corresponding EID-Prefixes through BGP to the DFZ. These PITRs will then encapsulate the traffic they receive for those EIDs towards the RLOCs of the LISP site, ensuring their reachability from non-LISP sites.

While it is possible for a PSP to manually configure each client's EID-Routes to be announced, this approach offers little flexibility and is not scalable. This section presents a scalable architecture that offers automatic distribution of EID-Routes to LISP sites and service providers.

The architecture requires no modification to existing LISP network elements, but it introduces a new (conceptual) network element, the EID-Route Server, which is defined as a router that either propagates routes learned from other EID-Route Servers or originates EID-Routes. The EID-Routes that it originates are those for which it is authoritative. It propagates these routes to Proxy-ITRs within the AS of the EID-Route Server. It is worth noting that a BGP-capable router can also be considered an EID-Route Server.

Further, an EID-Route is defined as a prefix originated via the Route Server of the MSP, which should be aggregated if the MSP has multiple customers inside a single large continuous prefix. This prefix is propagated to other PITRs both within the MSP and to other Pitr operators with which it peers. EID-Route Servers are operated by either the LISP site, MSPs, or PSPs and may be collocated with a Map-Server or Pitr, but they are functionally discrete entities. They distribute EID-Routes, using BGP, to other domains according to policies set by participants.




Figure 7: PITR-RD Architecture

The architecture described above decouples EID origination from route propagation, with the following benefits:

- o Can accurately represent business relationships between Pitr operators
- o Is more mapping system agnostic

- o Makes minor changes to PIR implementation; no changes to other components

In the example in Figure 7, we have a MSP providing services to the LISP site. The LISP site does not run BGP and gets an EID allocation directly from a RIR, or from the MSP, which may be a Local Internet Registry (LIR). Existing PI allocations can be migrated as well. The MSP ensures the presence of the prefix in the mapping system and runs an EID-Route Server to distribute it to PSPs. Since the LISP site does not run BGP, the prefix will be originated with the AS number of the MSP.

In the simple case depicted in Figure 7, the EID-Route of a LISP site will be originated by the Route Server and announced to the DFZ by the PSP's PIRs with AS path 64501 64500. From that point on, the usual BGP dynamics apply. This way, routes announced by the PIR are still originated by the authoritative Route Server. Note that the peering relationships between MSPs/PSPs and those in the underlying forwarding plane may not be congruent, making the AS path to a PIR shorter than it is in reality.

The non-LISP site will select the best path towards the EID-Prefix according to its local BGP policies. Since AS-path length is usually an important metric for selecting paths, careful placement of PIRs could significantly reduce path stretch between LISP and non-LISP sites.

The architecture allows for flexible policies between MSPs/PSPs. Consider the EID-Route Server networks as control plane overlays, facilitating the implementation of policies necessary to reflect the business relationships between participants. The results are then injected into the common underlying forwarding plane. For example, some MSPs/PSPs may agree to exchange EID-Prefixes and only announce them to each of their forwarding plane customers. Global reachability of an EID-Prefix depends on the MSP from which the LISP site buys service and is also subject to agreement between the above-mentioned parties.

In terms of impact on the DFZ, this architecture results in a slower routing table increase for new allocations, since traffic engineering will be done at the LISP level. For existing allocations migrating to LISP, the DFZ may decrease, since MSPs may be able to aggregate the prefixes announced.

Compared to LISP+BGP, this approach avoids DFZ bloat caused by prefix deaggregation for traffic engineering purposes, resulting in slower routing table increase in the case of new allocations and potential decrease for existing ones. Moreover, MSPs serving different clients with adjacent aggregatable prefixes may lead to additional decrease, but quantifying this decrease is subject to future research study.

The flexibility and scalability of this architecture do not come without a cost, however: A PSP operator has to establish either transit or peering relationships to improve its connectivity.

5.4. Migration Summary

Registering a domain name typically entails an annual fee that should cover the operating expenses for publishing the domain in the global DNS. This situation is similar for several other registration services. A LISP MSP client publishing an EID-Prefix in the LISP mapping system has the option of signing up for Pitr services as well, for an extra fee. These services may be offered by the MSP itself, but it is expected that specialized PSPs will do it. Clients that do not sign up will be responsible for getting non-LISP traffic to their EIDs (using the LISP+BGP scenario).

Additionally, Tier 1 ISPs have incentives to offer Pitr services to non-subscribers in strategic places just to attract more traffic from competitors and thus more revenue.

The following table presents the expected effects that the transition scenarios at various phases will have on the DFZ routing table size:

Phase	LISP+BGP	MSP Pitr	Pitr-RD
Early transition	no change	slower increase	slower increase
Late transition	may decrease	slower increase	slower increase
LISP Internet	considerable decrease		

It is expected that Pitr-RD will coexist with LISP+BGP during the migration, with the latter being more popular in the early transition phase. As the transition progresses and the MSP Pitr and Pitr-RD ecosystem gets more ubiquitous, LISP+BGP should become less attractive, slowing down the increase of the number of routes in the DFZ.

Note that throughout Section 5 we focused on the effects of LISP deployment on the DFZ routing table size. Other metrics may be impacted as well but to the best of our knowledge have not been measured as yet.

6. Security Considerations

All security implications of LISP deployments are to be discussed in separate documents. [LISP-THREATS] gives an overview of LISP threat models, including ETR operators attracting traffic by overclaiming an EID-Prefix (Section 4.4.3 of [LISP-THREATS]). Securing mapping lookups is discussed in [LISP-SEC].

7. Acknowledgements

Many thanks to Margaret Wasserman for her contribution to the IETF 76 presentation that kickstarted this work. The authors would also like to thank Damien Saucez, Luigi Iannone, Joel Halpern, Vince Fuller, Dino Farinacci, Terry Manderson, Noel Chiappa, Hannu Flinck, Paul Vinciguerra, Fred Templin, Brian Haberman, and everyone else who provided input.

8. References

8.1. Normative References

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

8.2. Informative References

- [CACHE] Jung, J., Sit, E., Balakrishnan, H., and R. Morris, "DNS performance and the effectiveness of caching", IEEE/ACM Transactions on Networking (TON), Volume 10, Issue 5, pages 589-603, October 2002.
- [DDT-ROOT] "Introduction to LISP DDT: DDT Root", March 2014, <<http://ddt-root.org/>>.
- [LISP-DDT] Fuller, V., Lewis, D., Ermagan, V., and A. Jain, "LISP Delegated Database Tree", Work in Progress, March 2013.

- [LISP-SEC] Maino, F., Ermagan, V., Cabellos-Aparicio, A., Saucez, D., and O. Bonaventure, "LISP-Security (LISP-SEC)", Work in Progress, October 2013.
- [LISP-THREATS]
Saucez, D., Iannone, L., and O. Bonaventure, "LISP Threats Analysis", Work in Progress, April 2014.
- [RFC4459] Savola, P., "MTU and Fragmentation Issues with In-the-Network Tunneling", RFC 4459, April 2006.
- [RFC4786] Abley, J. and K. Lindqvist, "Operation of Anycast Services", BCP 126, RFC 4786, December 2006.
- [RFC4984] Meyer, D., Zhang, L., and K. Fall, "Report from the IAB Workshop on Routing and Addressing", RFC 4984, September 2007.
- [RFC6834] Iannone, L., Saucez, D., and O. Bonaventure, "Locator/ID Separation Protocol (LISP) Map-Versioning", RFC 6834, January 2013.
- [RFC6835] Farinacci, D. and D. Meyer, "The Locator/ID Separation Protocol Internet Groper (LIG)", RFC 6835, 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.
- [RFC6887] Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", RFC 6887, April 2013.
- [TELCO96] Federal Communications Commission, "Telecommunications Act of 1996", 1996, <<http://transition.fcc.gov/telecom.html>>.

Appendix A. Step-by-Step Example BGP-to-LISP Migration Procedure

To help the operational community deploy LISP, this informative section offers a step-by-step guide for migrating a BGP-based Internet presence to a LISP site. It includes a pre-install/pre-turn-up checklist, and customer and provider activation procedures.

A.1. Customer Pre-Install and Pre-Turn-Up Checklist

1. Determine how many current physical service provider connections the customer has, and their existing bandwidth and traffic engineering requirements.

This information will determine the number of routing locators, and the priorities and weights that should be configured on the xTRs.

2. Make sure the customer router has LISP capabilities.

- * Check the OS version of the CE router. If LISP is an add-on, check to see if it is installed.

This information can be used to determine if the platform is appropriate to support LISP, in order to determine if a software and/or hardware upgrade is required.

- * Have the customer upgrade (if necessary, software and/or hardware) to be LISP capable.

3. Obtain the current running configuration of the CE router. A suggested LISP router configuration example can be customized to the customer's existing environment.

4. Verify MTU handling.

- * Request an increase in MTU to 1556 or more on service provider connections. Prior to the MTU change, verify the transmission of a 1500-byte packet from the PxTR to the RLOC with the Don't Fragment (DF) bit set.
- * Ensure that the customer is not filtering ICMP Unreachable or Time Exceeded messages on their firewall or router.

LISP, like any tunneling protocol, will increase the size of packets when the LISP header is appended. If increasing the MTU of the access links is not possible, care must be taken that ICMP is not being filtered in order to allow Path MTU Discovery to take place.

5. Validate member prefix allocation.

This step checks to see whether the prefix used by the customer is a direct (Provider-Independent) prefix or a prefix assigned by a physical service provider (Provider Aggregatable). If the prefixes are assigned by other service providers, then a Letter of Agreement is required to announce prefixes through the Proxy Service Provider.

6. Verify the member RLOCs and their reachability.

This step ensures that the RLOCs configured on the CE router are in fact reachable and working.

7. Prepare for cut-over.

- * If possible, have a host outside of all security and filtering policies connected to the console port of the edge router or switch.
- * Make sure the customer has access to the router in order to configure it.

A.2. Customer Activating LISP Service

1. The customer configures LISP on CE router(s) according to the configuration recommended by the service provider.

The LISP configuration consists of the EID-Prefix, the locators, and the weights and priorities of the mapping between the two values. In addition, the xTR must be configured with Map-Resolver(s), Map-Server(s), and the shared key for registering to Map-Server(s). If required, Proxy-ETR(s) may be configured as well.

In addition to the LISP configuration:

- * Ensure that the default routes(s) to next-hop external neighbors is included and RLOCs are present in the configuration.
 - * If two or more routers are used, ensure that all RLOCs are included in the LISP configuration on all routers.
 - * It will be necessary to redistribute the default route via IGP between the external routers.
2. When transition is ready, perform a soft shutdown on existing eBGP peer session(s).
 - * From the CE router, use the LISP Internet Groper (LIG) [RFC6835] to ensure that registration is successful.
 - * To verify LISP connectivity, find and ping LISP connected sites. If possible, find ping destinations that are not covered by a prefix in the global BGP routing system, because PITRs may deliver the packets even if LISP connectivity is not working. Traceroutes may help determine if this is the case.
 - * To verify connectivity to non-LISP sites, try accessing a landmark (e.g., a major Internet site) via a web browser.

A.3. Cut-Over Provider Preparation and Changes

1. Verify site configuration, and then verify active registration on Map-Server(s).
 - * Authentication key.
 - * EID-Prefix.
2. Add EID space to map-cache on proxies.
3. Add networks to BGP advertisement on proxies.
 - * Modify route-maps/policies on PxTRs.
 - * Modify route policies on core routers (if non-connected member).
 - * Modify ingress policies on core routers.
 - * Ensure route announcement in looking glass servers, RouteViews.
4. Perform traffic verification test.
 - * Ensure that MTU handling is as expected (PMTUD working).
 - * Ensure Proxy-ITR map-cache population.
 - * Ensure access from traceroute/ping servers around Internet.
 - * Use a looking glass to check for external visibility of registration via several Map-Resolvers.

Authors' Addresses

Lorand Jakab
Cisco Systems
170 Tasman Drive
San Jose, CA 95134
USA

EEmail: lojakab@cisco.com

Albert Cabellos-Aparicio
Technical University of Catalonia
C/Jordi Girona, s/n
BARCELONA 08034
Spain

EEmail: acabello@ac.upc.edu

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

EEmail: fcoras@ac.upc.edu

Jordi Domingo-Pascual
Technical University of Catalonia
C/Jordi Girona, s/n
BARCELONA 08034
Spain

EEmail: jordi.domingo@ac.upc.edu

Darrel Lewis
Cisco Systems
170 Tasman Drive
San Jose, CA 95134
USA

EEmail: darlewis@cisco.com