draft-lisp-introduction-07

Change log -06

Albert Cabellos, Damien Saucez

IETF 91 – LISP WG

General

- Typos
- Editorial polishing
- Impersonal form
- Updates on the references
- s/internetworking/interworking/g
- Clarify that the Mapping System is not always publicly accessible

Abstract

- Explain the purpose of the draft

```
This document describes the architecture of the Locator/ID Separation
Protocol (LISP)
   architecture, its main operational mechanisms as well as its design
   rationale. (LISP), making it easier to read the rest of the LISP
specifications and providing a basis for discussion about the details
of the LISP protocols.  This document is used for introductory
purposes, more details can be found in RFC6830, the protocol
specification.
```

1 Introduction

- Fully rewritten
- Focus on the contents of the draft as opposed to the scalability issues of the Internet
- Straight to the point introduction

see http://www.ietf.org/mail-archive/web/lisp/current/msg05532.html

2 Definition of Terms

- Added, pointing the appropriate RFCs

```
2.   Definition of Terms

   This document describes the LISP architecture and does not define or
   introduce any new term.  The reader is referred to
   [RFC6830],[RFC6831],[RFC6832],[RFC6833],[RFC6834],[RFC6835],
   [RFC6836],[RFC7215] for the LISP definition of terms.
```

3 LISP Architecture

- Clarification on the definition of ITR/ETR

With LISP, LISP sites (edge) and the core of the Internet are ~~inter-connected~~
interconnected by means of LISP-capable routers (e.g., border ~~routers).~~
routers) using tunnels. When **packets originated from a LISP site are**
flowing towards the core network, they ~~provide egress (from~~ **ingress into an encapsulated**
tunnel via an Ingress Tunnel Router (ITR). When packets flow from
the core ~~perspective)~~ **network** to a LISP ~~site~~ **site,** they ~~are called~~ **egress from an encapsulated**
tunnel to an Egress Tunnel ~~Routers (ETR), Ingress Tunnel Routers~~
~~(ITR) when they provide ingress, and~~ **Router (ETR). An** xTR ~~when they provide both.~~
~~ITRs~~ **is a router with can**
perform both ITR and ~~ETRs exchange~~ **ETR operations. In this context ITRs**
encapsulate packets ~~by encapsulating~~ **while ETRs decapsulate** them, hence LISP operates
as an overlay to the current Internet core.

3.3.1 LISP Encapsulation

- Overview of what's included in a LISP header

```
3. LISP header that may contain reachability information
LISP header that contains various forwarding-plane features
(such as reachability) and an Instance ID field. This
header is originated by ITRs and stripped by ETRs.
```

3.4.1 LISP Mappings

- Clarification for LCAF

Typical mappings in LISP bind EIDs in the form of IP prefixes with a set of RLOCs, also in the form of IPs. ~~Such~~ **IPv4 and IPv6** addresses are encoded using ~~a~~ **the appropriate Address Family Identifier (AFI) [RFC3232]. However LISP can also support more** general ~~syntax called~~ **address encoding by means of the ongoing effort around the** LISP Canonical Address Format ~~(LCAF), specified in~~ **(LCAF)** [I-D.ietf-lisp-lcaf]. ~~The syntax is general enough to support encoding of IPv4 and IPv6 addresses and any other type of value.~~

3.4.3.1 LISP+ALT

- Correction and clarification on how ALT works

The LISP Alternative Topology (LISP+ALT) [RFC6836] was the first Mapping System proposed, developed and deployed on the LISP pilot network. It is based on a distributed BGP ~~overlay. All the participating~~ **overlay participated by Map-Servers and Map-Resolvers. The** nodes connect to their peers through static tunnels. ~~Every ETR~~ **Each Map-Server** involved in the ALT topology advertises ~~its EID prefixes~~ **the EID-prefixes registered by the serviced ETRs,** making the EID routable on the ~~overlay.~~ **ALT topology.**

When an ITR needs a ~~mapping,~~ **mapping** it sends a Map-Request to a ~~nearby ALT router. The ALT routers then forward the Map-Request on~~ **Map-Resolver that, using** the ~~overlay by inspecting their~~ ALT ~~routing tables. When~~ **topology, forwards** the Map-Request ~~reaches~~ **towards** the ~~ETR~~ **Map-Server** responsible for the ~~mapping, a Map-Reply is generated and~~ **mapping. Upon reception the Map-Server forwards the request to the ETR that in turn, replies** directly ~~sent~~ to the ~~ITR's RLOC, without~~ **ITR** using the ~~ALT overlay.~~

~~2.4.3.2.~~ **native Internet core.**

3.4.3.2 LISP-DDT

- Clarification on the role of MR in DDT

```
LISP-DDT [I-D.ietf-lisp-ddt] is conceptually similar to the
DNS, a hierarchical directory whose internal structure
mirrors the hierarchical nature of the EID address space.
The DDT hierarchy is composed of DDT nodes forming a tree
structure, the leafs of the tree are Map-Servers. On top of
the structure there is the DDT root node [DDT-ROOT], which
is a particular instance of a DDT node and that matches the
entire address space. As in the case of DNS, DDT supports
multiple redundant DDT nodes and/or DDT roots. The
following figure presents a schematic representation
Finally, Map-Resolvers are the clients of the DDT
hierarchy. hierarchy and can query either the DDT root and/
or other DDT nodes.
```

4.1 Cache Management

- Additional information on how the Map-Cache can operate (proactive caching)

```
Finally it is worth noting that in some cases an entry in
the map-cache can be proactively refreshed using the
mechanisms described in the section below.
```

4.2 RLOC Reachability

- Additional information on the joint operation of RLOC probing and Echo-nonce

```
It is worth noting that RLOC probing and Echo-nonce can work
  together.  Specifically if a nonce is not echoed, an ITR could RLOC-
  probe to determine if the path is up because the return bidirectional
  path may have failed or the return path is not used, that is there is
  only a unidirectional path.
```

4.4 MTU Handling

- Clarification on both stateless and stateful MTU handling

Since LISP encapsulates packets it requires dealing with packets that
exceed the MTU of the path between the ITR and the ETR. Specifically
LISP ~~defienes~~ **defines** two mechanisms:

Stateless: With this mechanism ~~ITRs fragment packets that are~~ **the effective MTU is assumed from the ITR's perspective. If a payload packet is** too
~~big, typically~~ **big for the effective MTU, and can be fragmented, the payload packet is fragmented on the ITR, such that** reassembly is performed at the
destination host.

Stateful: With this mechanism ITRs keep track of the MTU of the
paths towards the destination locators by parsing the ICMP Too Big
packets sent by intermediate routers. **Additionally ITRs will send ICMP Too Big messages to inform the sources about the effective MTU.**

In both cases if the ~~packet cannot~~ **packet cannot be fragmented (IPv4 with DF=1 or IPv6) then the ITR drops it and replies with a ICMP Too Big message to the source.**

5. Mobility

- Clarification on LISP mobility

The separation between locators and identifiers in LISP was initially proposed for traffic engineering purpose where LISP sites can change their attachment points to the Internet (i.e., RLOCs) without impacting endpoints or the Internet core. In this context, the border routers operate the xTR functionality and endpoints are not aware of the existence of LISP. However, this mode of operation does not allow seamless mobility of endpoints between different LISP sites as the EID address might not be ~~framgneted (IPv4 with DF=1 or IPv6) then the ITR drops it and replies with~~ routable in a ~~ICMP Too Big message to the source.~~

~~4. Mobility~~ visited site.

Nevertheless, LISP can ~~also~~ be used to enable seamless IP mobility ~~of devices not located in~~ when LISP ~~networks. The problem with mobility of such devices~~ is ~~that their IP address changes whenever they change location, interrupting so flows.~~

~~To enable mobility on such devices, the device can implement~~ directly implemented in the endpoint. Each endpoint is then an xTR ~~functionality where~~ and the ~~IP~~ EID address ~~presented to applications~~ is ~~an EID that never changes while~~ the ~~IP address obtained from~~ one presented to the network ~~is~~ stack used by applications while the ~~xTR as RLOC. Packets are then transported on the network using~~ RLOC is the ~~IP~~ address ~~assigned to the device by~~ gathered from the ~~visited~~ network ~~while at the application level IP addresses remain independent of the location of the device.~~ when it is visited.

6. Multicast

- Correction on IGMP

    1. An end-host that belongs willing to join a multicast channel sends an IGMP report. Multicast PIM routers at the LISP site transmits a propagate PIM Join/ Prune message (S-EID,G) to join a multicast group. Join/Prune messages (S-EID, G) towards the ETR.

- Additional information on non-PIM multicast

    LISP also support non-PIM mechanisms to maintain multicast state.

8.3 LISP for Virtual Private Networks

- Remove the notion to BGP as it is not the only way of doing VPNs

```
It is ~~nowadays~~ common to operate several virtual networks over the same
physical infrastructure.  ~~The current approach usually rely on
BGP/MPLS VPNs, where BGP is used to exchange routing information and
MPLS to segregate packets of the different logical networks.  This
functionality could be achieved with LISP where the mappings and the
mapping system are used instead of BGP and the LISP encapsulation is
used to replace MPLS.~~  In **such** virtual **private** networks, it is
essential to distinguish to which virtual network a packet belongs
and tags or labels are used for that purpose.  With LISP, the
distinction can be made with the Instance ID field.  When an ITR
encapsulates a packet from a particular virtual network (e.g., known
via the VRF or VLAN), it tags the encapsulated packet with the
Instance ID corresponding to the virtual network of the packet.  When
an ETR receives a packet tagged with an Instance ID it uses the
Instance ID to determine how to ~~threat~~ **treat** the packet.

~~Appart from the simplicity of managing mappings, the~~

**The main** advantage of using LISP for virtual ~~network~~ **networks, on top of the
simplicity of managing the mappings,** is that it does not impose any
requirement on the underlying network, ~~except~~ **as long as it is** running IP.
```

8.4. LISP for Virtual Machine Mobility in Data Centers

- Removed because incorrect and simplified for clarity

```
   To inform the other LISP routers that the machine moved and where,
   and then to avoid detours via the initial subnetwork, every Map-
   Server can listen on a predefined multicast address that is used mechanisms such
   as
   source address for Map-Register.  As a result, the Map-Notify sent
   back by the Map-Server will be received by all the LISP routers that
   hence automatically learn the new location of the virtual machine.

8. Solicit-Map-Request messages are used.
```

10. Acknowledgements

- Added

11. **Acknowledgements**

This document was initiated by Noel Chiappa and much of the core philosophy came from him. The authors acknowledge the important contributions he has made to this work and thank him for his past efforts.

The authors would also like to thank Dino Farinacci, Fabio Maino, Luigi Iannone, Sharon Barakai, Isidoros Kouvelas, Christian Cassar, Florin Coras, Marc Binderberger, Alberto Rodriguez-Natal, Ronald Bonica, Chad Hintz, Robert Raszuk, Joel M. Halpern, Darrel Lewis, as well as every people acknowledged in [RFC6830].

Appendix A. A Brief History of Location/Identity Separation

- Correction

```
A small group of like-minded personnel from various scattered
locations within Cisco, spontaneously formed immediately after that
workshop, to work on an idea that came out of informal discussions at
the ~~workshop.~~ **workshop and on various mailing lists.**  The first Internet-Draft
on LISP appeared in January,
~~2007, along with a LISP mailing list at the IETF.~~ **2007.**
```

Change log -07

General

- A few typos
- Some minor corrections (words, etc)

3.1 Design Principles

- Clarification on the benefits of decoupled and control-plane separation

```
Decoupled data and control-plane: Separating the data-plane from the
control-plane allows them to scale independently and use different
architectural approaches. This is important given that they
typically have different requirements. requirements and allows for
other data-planes to be added.
```

3.4.2 Mapping System Interface

- Additional information (Map-Notify)

```
Map-Notify: When requested by the ETR, this message is sent by the Map-
Server in response to a Map-Register to acknowledge the correct
reception of the mapping and convey the latest Map-Server state on the
EID to RLOC mapping. In some cases a Map-Notify can be sent to the
previous RLOCs when an EID is registered by a new set of RLOCs.
```

8.4 LISP for VM Mobility in DC

- Clarification through simplification

A way to enable seamless virtual machine mobility in data center is to conceive the datacenter backbone as the RLOC space and the subnet where servers are hosted as forming the EID space. A LISP router is placed at the border between the backbone and each subnet. When a virtual machine is moved to another subnet, it can (temporarily) keep (temporarily) the address of the subnet it was hosted had before the move so to allow ongoing communications to subsist. When continue without a subnet transport layer connection reset. When an xTR detects the presence of a host with source address received on a subnet to be an address that does not belong assigned to the subnet (e.g., via a message sent by the hypervisor or traffic inspection), the LISP router of the new subnet subnet, it registers the IP address of the virtual machine as an EID to the Map-Server of the subnet and associates its own address as RLOC. Mapping System.

Change log -08

Change log -08

- Few editorial editorial changes for clarification and to address Ron's comments.