

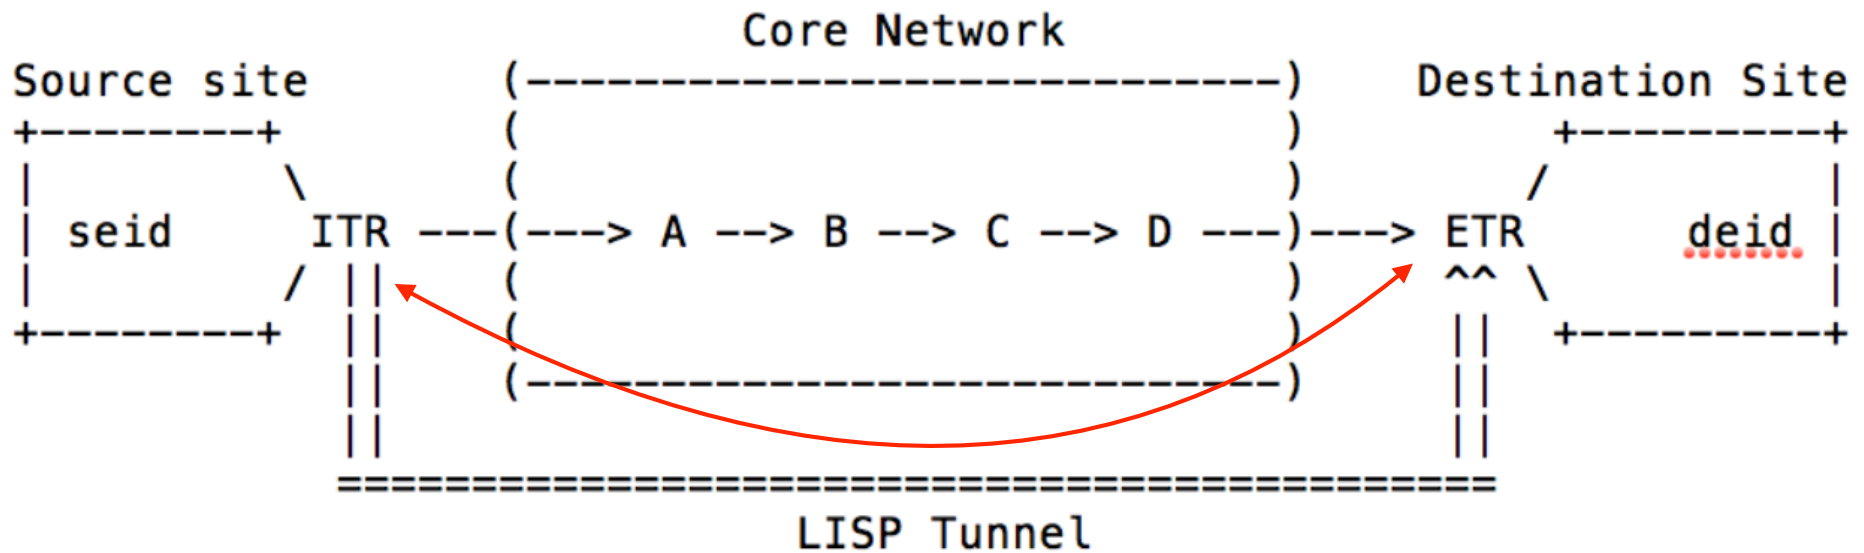
LISP Traffic Engineering

`draft-farinacci-lisp-te-12`

*LISP Working Group - Chicago IETF
March 2017*

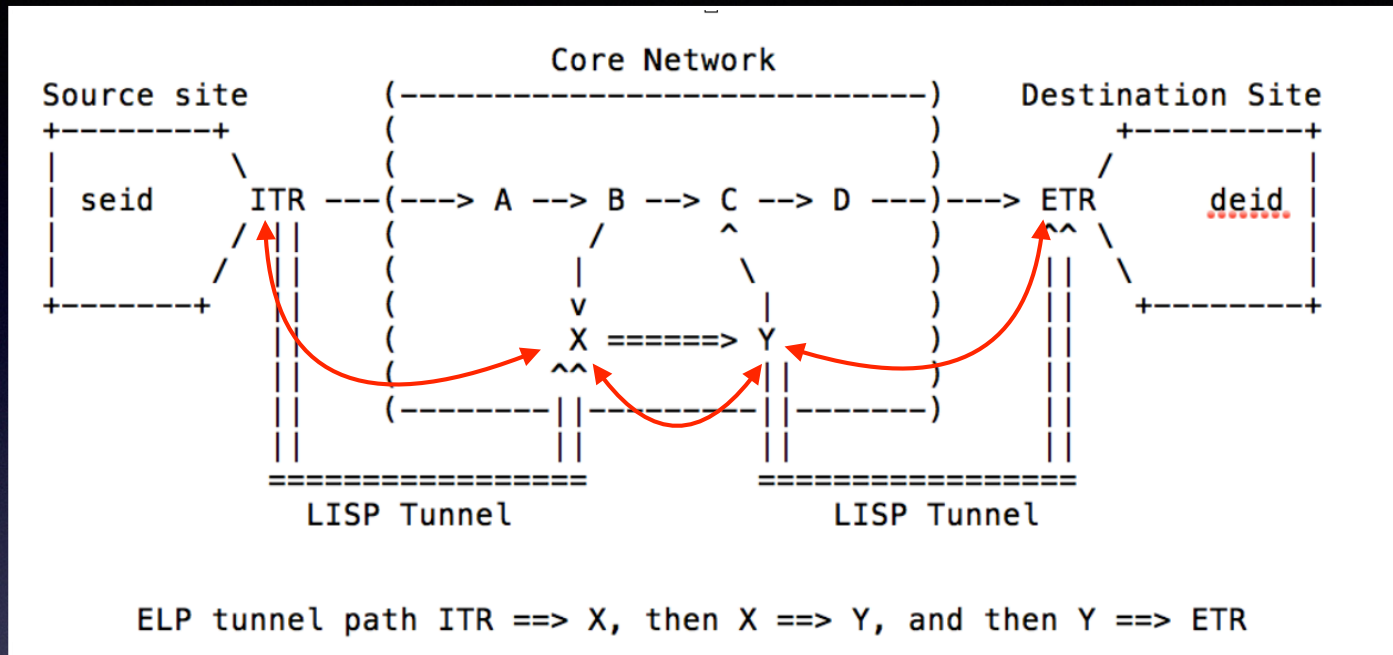
Dino Farinacci, Michael Kowal, Parantap Lihiri

Typical Overlay



Typical Data Path from ITR to ETR

LISP-TE



There are various reasons why the path from 'seid' to 'deid' may want to avoid the path from B to C. To list a few:

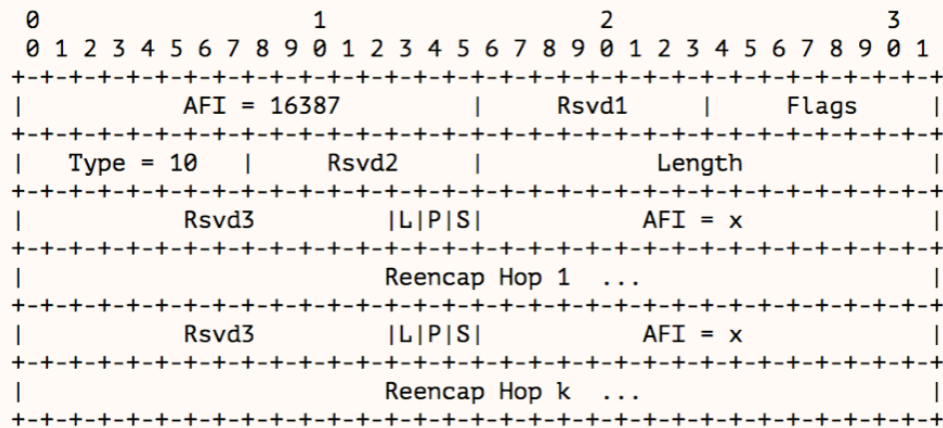
- o There may not be sufficient capacity provided by the networks that connect B and C together.
- o There may be a policy reason to avoid the ASes that make up the path between B and C.
- o There may be a failure on the path between B and C which makes the path unreliable.
- o There may be monitoring or traffic inspection resources close to RTRs X and Y that do network accounting or measurement.
- o There may be a chain of services performed at RTRs X and Y regardless if the path from ITR to ETR is through B and C.

ELP Use-Case

4.6. Traffic Engineering Using Re-encapsulating Tunnels

For a given EID lookup into the mapping database, this LCAF can be returned to provide a list of locators in an explicit re-encapsulation path. See [LISP-TE] for details.

Explicit Locator Path (ELP) Canonical Address Format:



Length: length in bytes starting and including the byte after this Length field.

Rsvd3: this field is reserved for future use and MUST be transmitted as 0 and ignored on receipt.

Lookup bit (L): this is the Lookup bit used to indicate to the user of the ELP not to use this address for encapsulation but to look it up in the mapping database system to obtain an encapsulating RLOC address.

RLOC Probe bit (P): this is the RLOC Probe bit that means the Reencap Hop allows RLOC-probe messages to be sent to it. When the R bit is set to 0, RLOC-probes must not be sent. When a Reencap Hop is an anycast address then multiple physical Reencap Hops are using the same RLOC address. In this case, RLOC-probes are not needed because when the closest RLOC address is not reachable, another RLOC address can be reachable.

Strict bit (S): this is the Strict bit, which means the associated Reencap Hop is required to be used. If this bit is 0, the re-encapsulator can skip this Reencap Hop and go to the next one in the list.

AFI = x: x can be any AFI value from [AFN]. When a specific AFI has its own encoding of a multicast address, this field must be either a group address or a broadcast address.

Usage: This encoding can be used in RLOC-records in Map-Request, Map-Reply, Map-Register, and Map-Notify messages. This encoding does not need to be understood by the mapping system for mapping database lookups, since this LCAF Type is not a lookup key.

RFC8060 (LCAF)

Different Paths to RLOCs

```
EID-prefix: 10.0.0.0/8
Locator-set: ETR-A: priority 1, weight 50
              ETR-B: priority 1, weight 50
              ETR-C: priority 2, weight 50
              ETR-D: priority 2, weight 50
```

If an ELP is going to be used to have a policy path to ETR-A and possibly another policy path to ETR-B, the locator-set would be encoded as follows:

```
EID-prefix: 10.0.0.0/8
Locator-set: (x, y, ETR-A): priority 1, weight 50
              (q, r, ETR-B): priority 1, weight 50
              ETR-C:         priority 2, weight 50
              ETR-D:         priority 2, weight 50
```

Source-Based Forwarding

Using source/destination lookups into the mapping database can yield different ELPs. So for example, a premium service flow with (source=1.1.1.1, dest=10.1.1.1) can be described by using the following mapping entry:

```
EID-prefix:    (1.0.0.0/8, 10.0.0.0/8)
Locator-set:   (x, y, ETR-A): priority 1, weight 50
               (q, r, ETR-A): priority 1, weight 50
```

And all other best-effort sources would use different mapping entry described by:

```
EID-prefix:    (0.0.0.0/0, 10.0.0.0/8)
Locator-set:   (x, x', y, y', ETR-A): priority 1, weight 50
               (q, q', r, r', ETR-A): priority 1, weight 50
```


Multicast Forwarding

```
Multicast EID key: (seid, G)
Locator-set:      ETR-A: priority 1, weight 25
                  ETR-B: priority 1, weight 25
                  g1:   priority 1, weight 25
                  g2:   priority 1, weight 25
```

```
Multicast EID key: (seid, G)
Locator-set:      (x, y, ETR-A): priority 1, weight 50
                  (a, g, b, ETR-B): priority 1, weight 50
```

Using ELPs for multicast flows

Document History



No changes
for **4 years!**

Content Update - Jul 2012

LISP-TE initial draft
Mar 2012

- o Add the Lookup bit to allow an ELP to be a list of encapsulation and/or mapping database lookup addresses.
- o Indicate that ELPs can be used for service chaining.
- o Add text to indicate that Map-Notify messages can be sent to new RTRs in a ELP so their map-caches can be pre-populated to avoid mapping database lookup packet loss.
- o Fixes to editorial comments from Gregg.

Chicago IETF
Mar 2017

Request for Working Group Document

