


# LISP Uberlay

draft-moreno-lisp-uberlay

Dino Farinacci,  
Sanjay Hooda,  
Satish Kondalam  
Fabio Maino,  
Victor Moreno,  
Marc Portoles Comeras,  
Alberto Rodriguez-Natal

IETF 103 – Bangkok, Thailand  
November, 2018

# Uberlay: Interconnection of multiple overlays



- Structure the LISP network hierarchically
  - Many site-overlays interconnected by a single transit “uberlay”
  - Only one uberlay, no backdoors
- Provide fate-isolation & site-overlay survivability
- Assist mapping system scale
  - Horizontally distribute state across site-overlays
  - Limit state to local EID-prefixes
- Segment RLOC space: private addressing, improved scale.

# Multi-overlay Control Plane

- Site-overlay mappings are registered to local mapping system
- Border xTRs register 'default' mapping with local site **RLOCs**
- Border xTRs subscribe to all local mappings (0/0)
- Border xTRs register site-overlay **EIDs** with 'uberlay' **RLOCs** in Uberlay mapping system
- Site-overlay routes on site-overlay **RLOCs** only
- Site-overlay mapping system - not necessary to run LISP-DDT


- Uberlay has its own mapping system to collect site-overlay **EIDs** with uberlay **RLOCs**
- Uberlay underlay routes on uberlay **RLOCs** only
- Uberlay mapping system may run LISP-DDT

## Benefits:

- Site-overlay mapping system has less state
- Site-overlay **RLOCs** not in Uberlay underlay
- Uberlay **RLOCs** only for borders are in site-overlay underlay
- Site-overlay **RLOCs** are private to site
- Reduced lookup latency for local **EIDs**
- Reduced registration latency for local **EIDs**

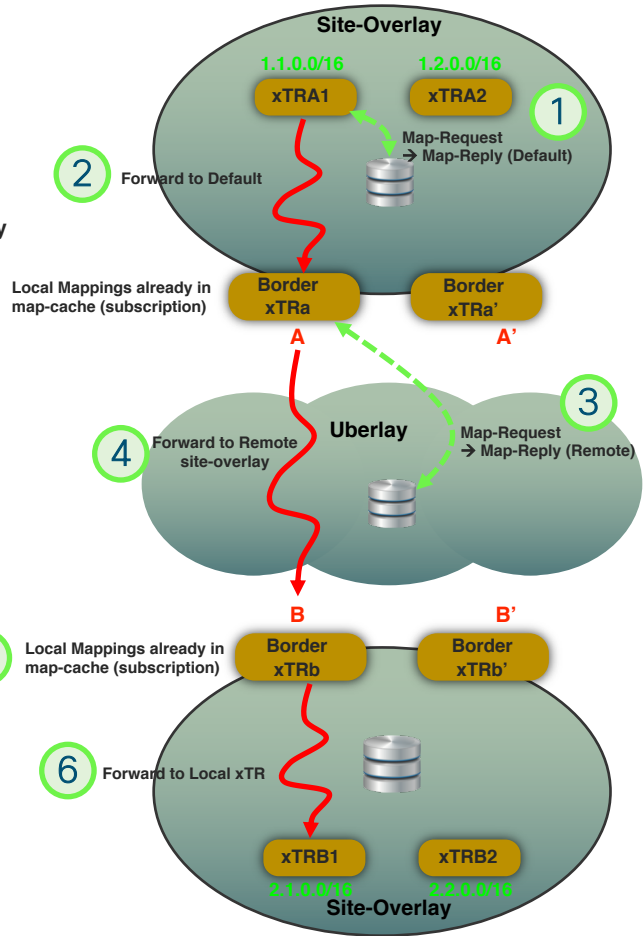
## Drawbacks:

- Uberlay mapping system stores site-overlay **EIDs** but can be aggregated (modulo mobility)



# Multi-overlay Forwarding

- Border xTRs subscribe to all local mappings (0/0)
- Site-overlay mappings are in the map-cache of the border xTRs
- xTR tunnels traffic to 'default' (border xTR) when destination EID is a remote site-overlay EID
- xTR de-encapsulates received traffic
- If destination EID not found in map-cache, query the uberlay Mapping System
- Tunnel traffic to remote border xTR based on Map-reply
- Lookup map-cache at border xTR, if destination EID not found, query the uberlay MS
- Tunnel traffic to remote border xTR based on Map-reply



**Site Overlay**  
 1.1.0.0/16 -> xTRA1  
 1.2.0.0/16 -> xTRA2  
 'default' -> (xTRa, xTRa')

**Uber Overlay**  
 1.1.0.0/16 -> (A, A')  
 1.2.0.0/16 -> (A, A')  
 2.1.0.0/16 -> (B, B')  
 2.2.0.0/16 -> (B, B')

**Site Overlay**  
 2.1.0.0/16 -> xTRB1  
 2.2.0.0/16 -> xTRB2  
 'default' -> (xTRb, xTRb')

# Border xTR role

- Connect site-overlays to the uberlay
  - Re-encapsulating Tunnel Routers (RTRs)
  - Exchange EID info between site-overlay & uberlay mapping systems
    - Constrain advertisements into site-overlay
    - Split-horizon for Uberlay learnt EIDs (mobility)
  - Register default mapping into site-overlay
- Interconnect separate RLOC spaces:
  - Site-overlay facing RLOC-set
  - Uberlay facing RLOC-set
- EID mobility:
  - Site-overlay facing and uberlay facing away table
  - Relay mobility events between site-overlays and uberlay

# Inter-site-overlay EID Mobility

- EID mobility in each site-overlay and in the uberlay follows the procedures in draft-ietf-lisp-eid-mobility. The following additions enable Inter-site-overlay mobility:
- The border xTRs on path will receive the Map-notify messages part of the eid-mobility procedures
  - Install the roaming EID in the away table facing the arrival overlay
  - Registration and eid-mobility procedures are triggered in the next overlay

# Inter-site Multicast

- The procedures in RFC8378 are followed in each domain
- A Map-notify is sent towards the RP or multicast source by the mapping system in each overlay
  - Border-xTRs are the on path and receive this map-notify
- The Map Notify in one overlay triggers the registration of multicast interest in the next overlay towards the source.
- Multicast replication lists/trees are built in each overlay
- The border xTRs provide a natural point of replication

# VPN Considerations

- Two sets of instance IDs:
  - Reuse set: Site-overlay specific, can be reused across site-overlays. Provide intra-site-overlay connectivity only.
  - Global set: Used across site-overlays and the uberlay. Provide inter-site-overlay connectivity.



# Next steps

- Discuss first draft in the mailing list
- Improve state reduction in uberlay
- Deeper exploration of VPN implications

Comments, Questions