

LISP Uberlay

draft-moreno-lisp-uberlay


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

IETF 106 – Singapore
November, 2019

Update at IETF 106

- Based on feedback from the WG and Chairs, we are updating the text with:
 - Discussion on topology considerations (loop free topology)
 - Considerations on resiliency of borders
 - Provisioning of a default EID
- Editorial updates have been made to other parts of the text
- Revision -02 will be posted soon

Uberlay: Recap on idea and motivation



- Provide interoperability between disparate site-overlay implementations
 - Control Plane: Different models at each site-overlay: DDT, Decent, ALT or other
 - Data Plane: Different encapsulations at each site-overlay: LISP, VXLAN, GPE, Geneve, other
- Structure the LISP network hierarchically
 - Many site-overlays interconnected by a transit “uberlay”
- Different RLOC spaces: private addressing, improved scale.
- Provide fate-isolation & site-overlay survivability

Updates in this version
(-02)

Topology support

- How do we stitch together multiple networks using overlays?
- Requirement: Loop free topologies

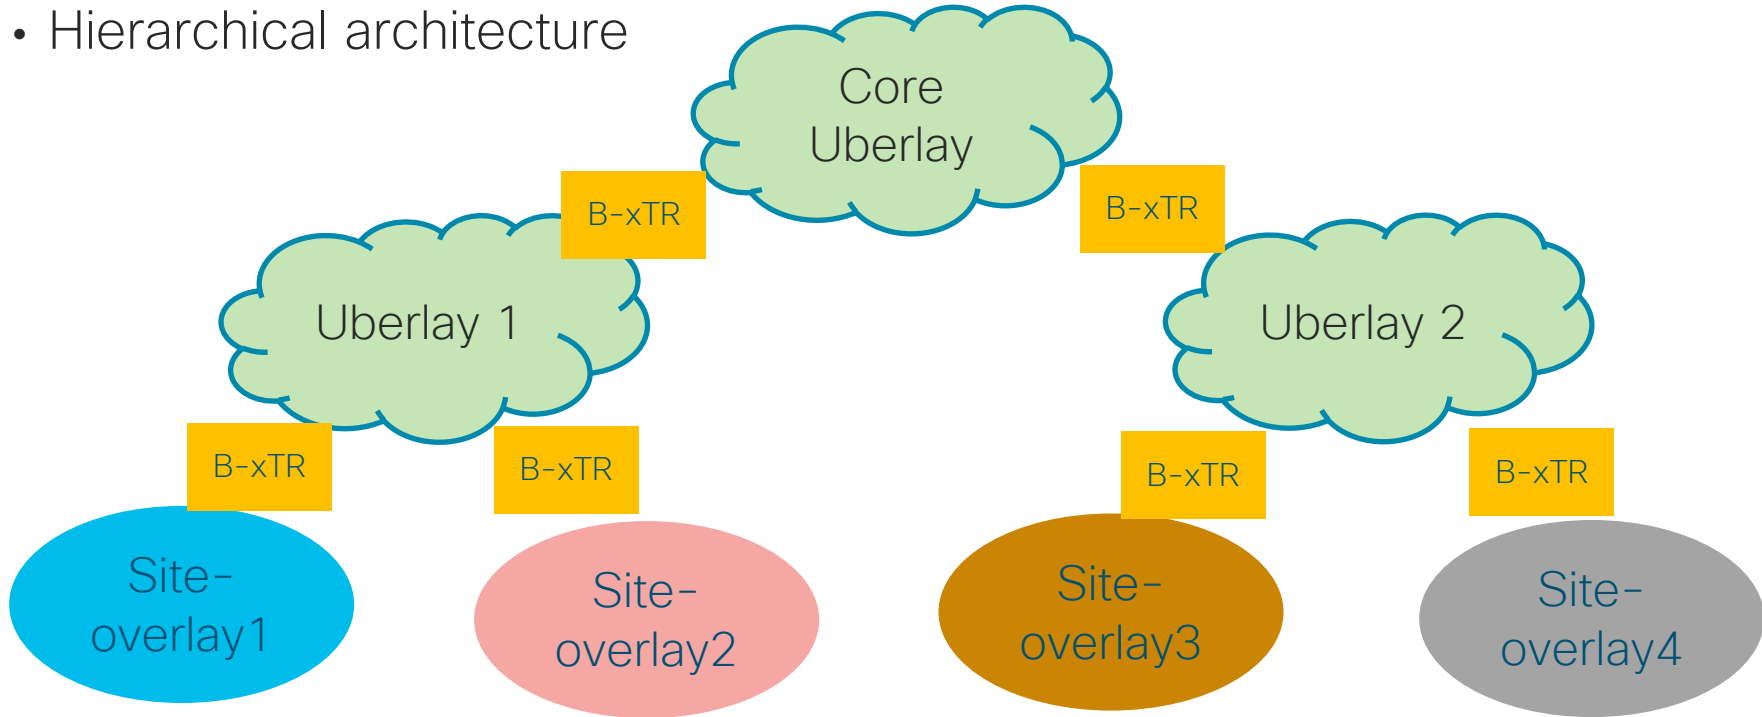
<< There is no provision for the detection of registration loops when concatenating site-overlays and Uberlays, thus any interconnection of overlay domains (site-overlays or Uberlays) must be done in a *loop free topology*.>>

Topology support discussion

- Loop free topologies
- General concept not encoded into protocol messages
- Limits peering of overlays to follow strict hierarchy
- Defined by two rules:
 - Uberlays must only connect to Uberlays in the next consecutive level of hierarchy
 - Uberlays within the same level of hierarchy must not connect to each other.

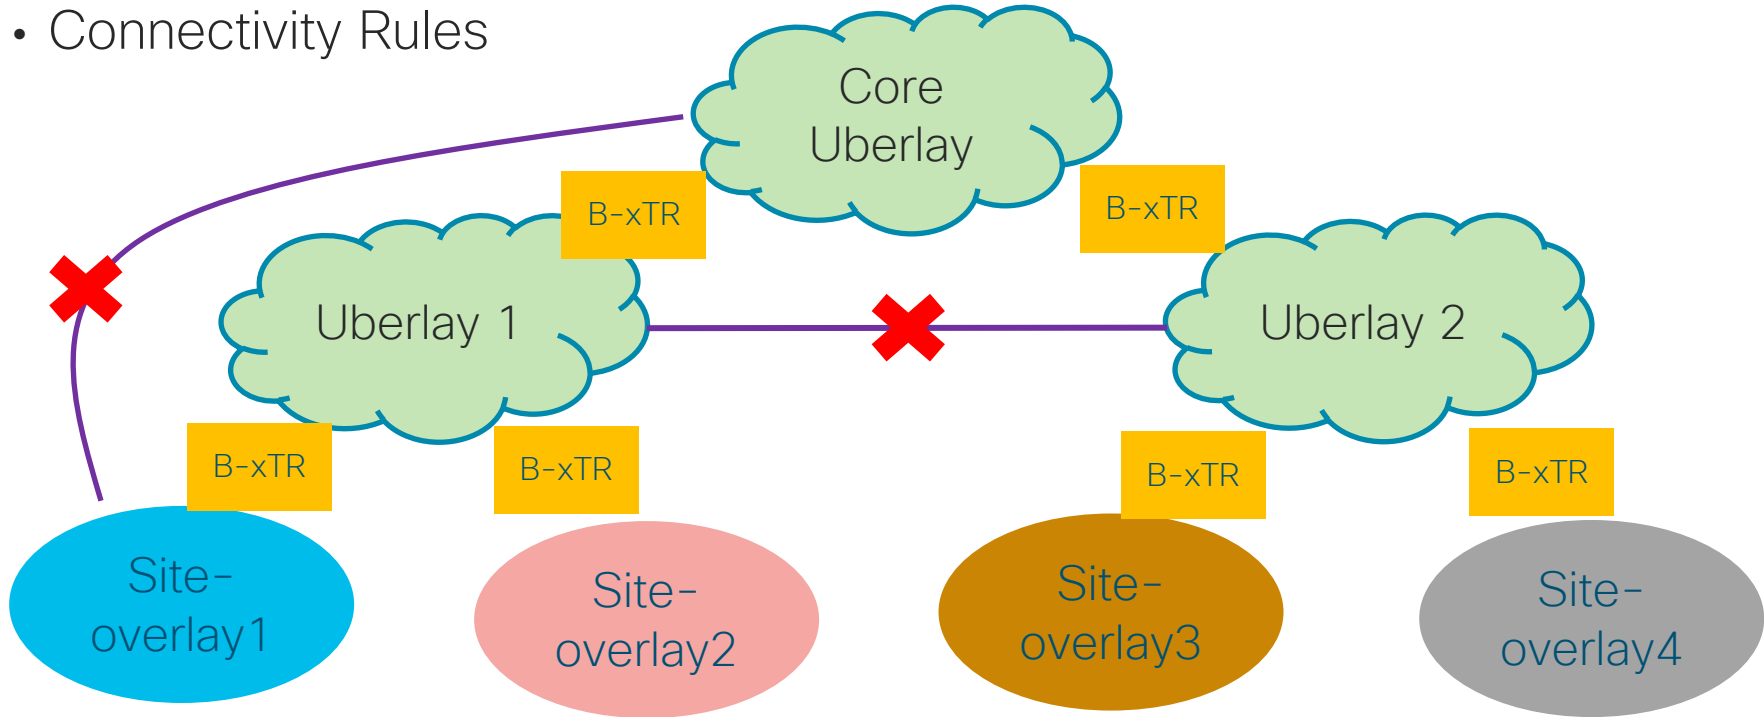
Topology support

- Hierarchical architecture



Topology support discussion

- Connectivity Rules



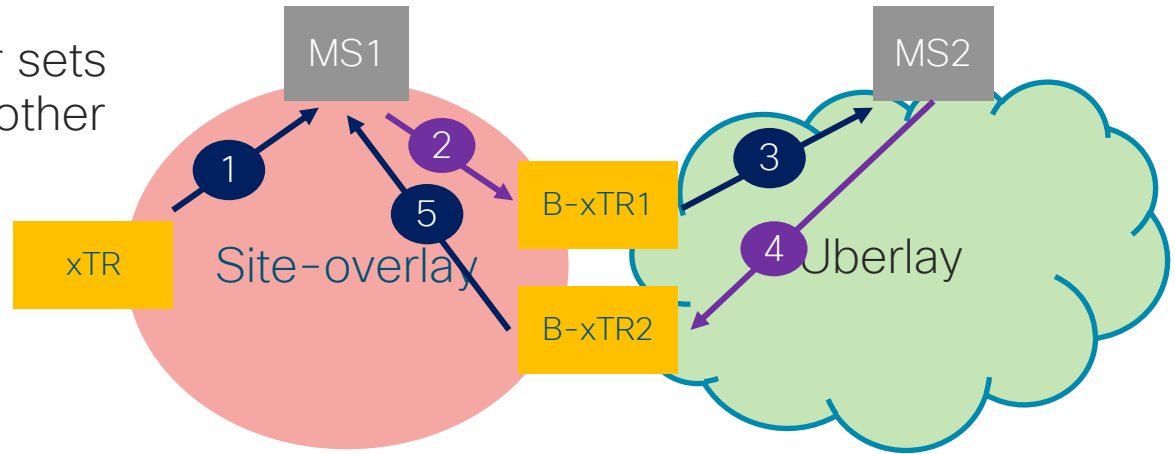
Border redundancy/resiliency support

- Requirement: multihoming and avoid registration loops

<< Redundancy at the border xTRs requires that border xTRs be logically grouped so that the redundant array doesn't create a registration loop. >>

Border redundancy/resiliency support

- Registration Loop
- Avoid loops:
 - Border xTRs in Border sets need to identify each other
 - Split horizon of registrations at Border set level



Border redundancy/resiliency support

- Border-sets and site-of-origin

<< Members of a border set will register the EIDs from a particular site-of-origin into the neighboring overlay (site-overlay or uberlay) using a common site-id>>

- Site-ID
 - Identifies common site-of-origin for a border-set
 - Locally significant to each overlay site
 - site-ID included in registration and publish messages
 - Border xTR does not register EIDs back to site-of-origin based on site-ID
 - Can be used to identify LISP and non-LISP sites-of-origin

Default EID registrations and treatment

- Default EID registrations

<<Border xTRs will register a mapping to be used as a default mapping to handle the forwarding of traffic destined to any EIDs that are not explicitly registered>>

<<This registration is intended to instruct the Mapping System to follow the procedures in [RFC6833] for **Negative Map Replies** (...) and issue a **map-reply** with the calculated EID and the RLOCs registered by the border xTRs>>

Comments, Questions