



IETF 101 – London
Mar 2018
SPRING Working Group

draft-ali-spring-srv6-oam-00.txt

SRv6 OAM

Z. Ali, C. Filsfils, et al (Cisco Systems)
J. Leddy (Comcast)
S. Matsushim (Softbank)
R. Raszuk (Bloomberg LP)
B. Peirens (Proximus)
G. Naik (Drexel University)

Reference Topology



k SRv6 Capable

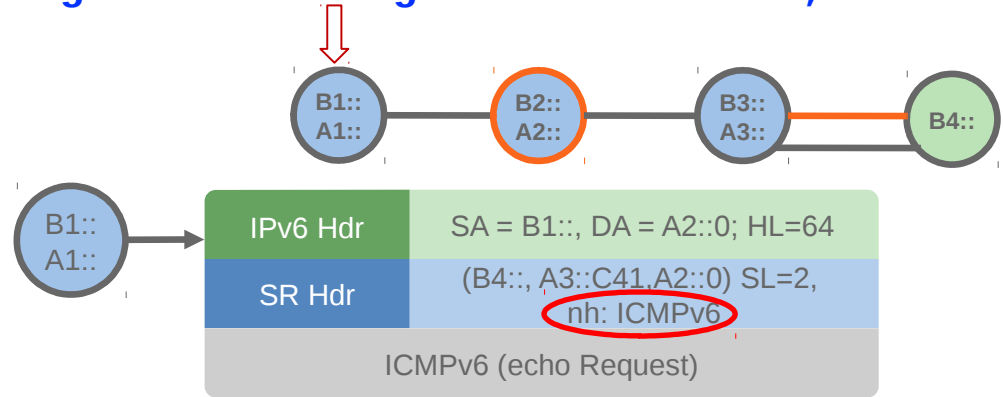
k classic IPv6 Node

k classic IPv4 Node

↓
○ Message Processing Node

ICMPv6 Ping Via a SID list


ping B4::
via segment-list A2::
A3::
C41



- The originator node constructs an SRH using the segment list specified by the user and adds it to IPv6 packet.
- All other ICMPv6 related processing remains unchanged.
- No changes are required at the transit node.
- No changes are required at the destination node.

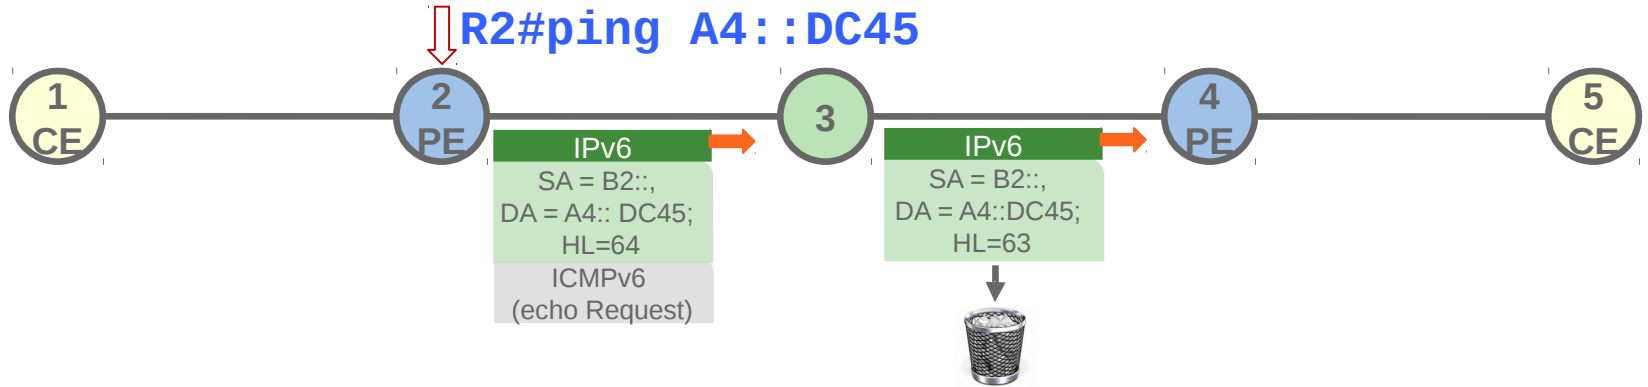
Traceroute Via a SID list

```
# traceroute B4:: via  
segment-list A2::0,A3::C41
```



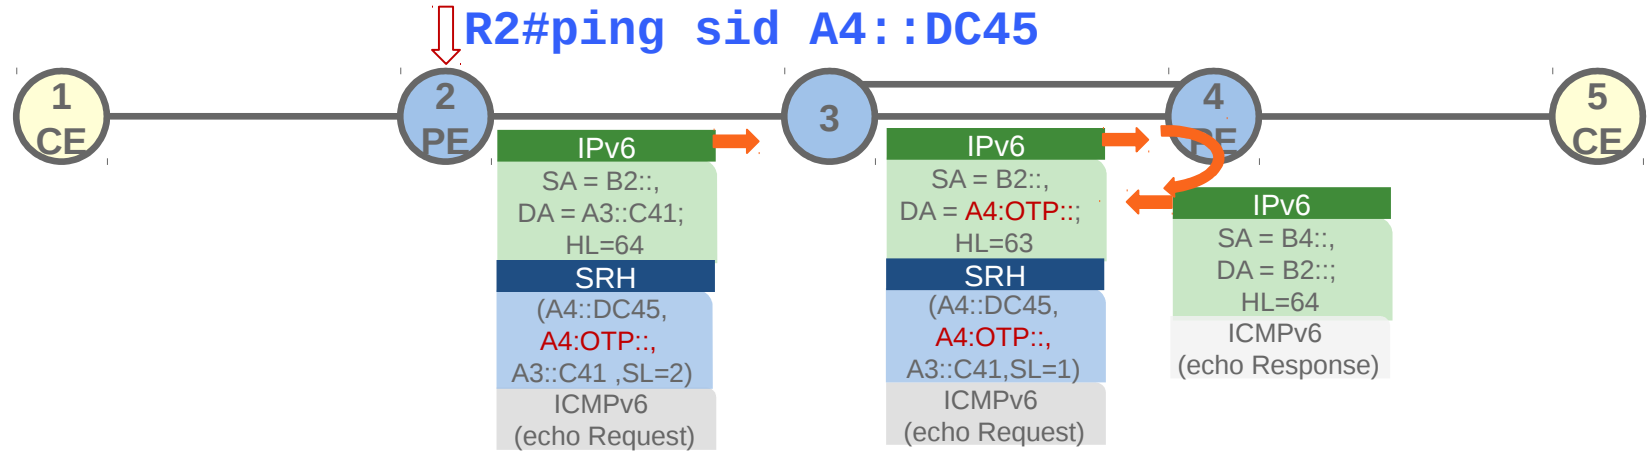
- The originator node constructs an SRH using the segment list specified by the user and adds it to traceroute probe packet.
- All other IPv6 traceroute related processing remains unchanged.
- No changes are required at the transit node.
- No changes are required at the destination node.

ICMPv6 Ping to a SID Function



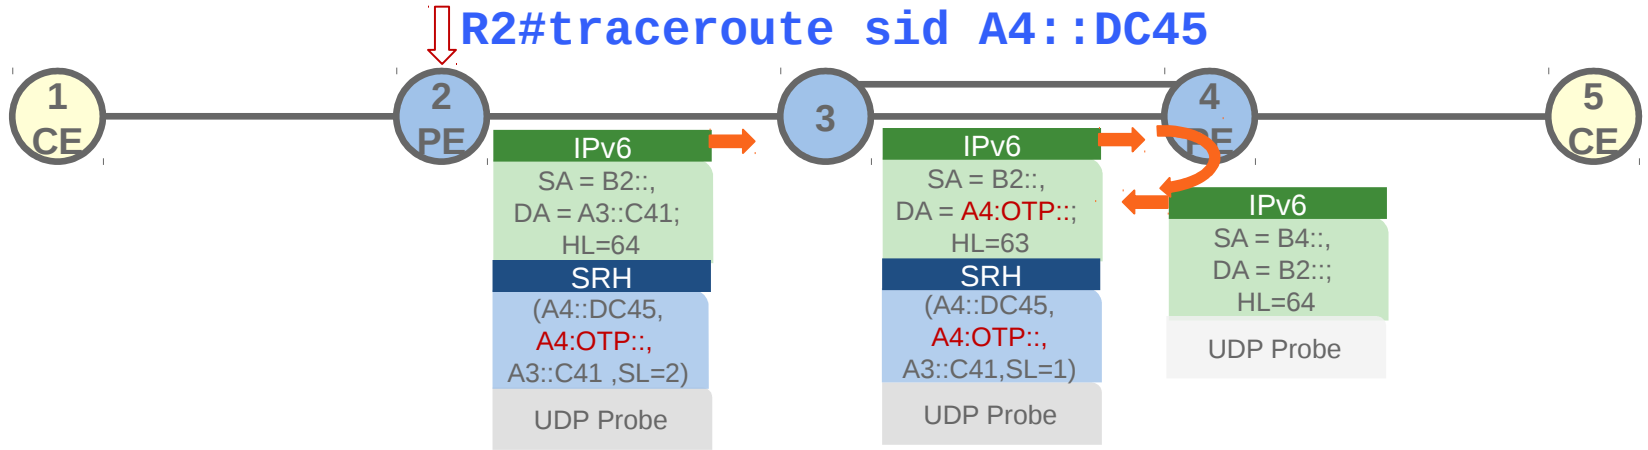
- This is the case of the use of IPv6 ping to ping a SID function.
- Forwarding chain on Egress will be incomplete; Egress cannot forward a packet with DA = A4::DC45 (e.g., there is no inner IP header).
- Ping to a SID function requires punting for OAM packets.

Pinging a SID Function



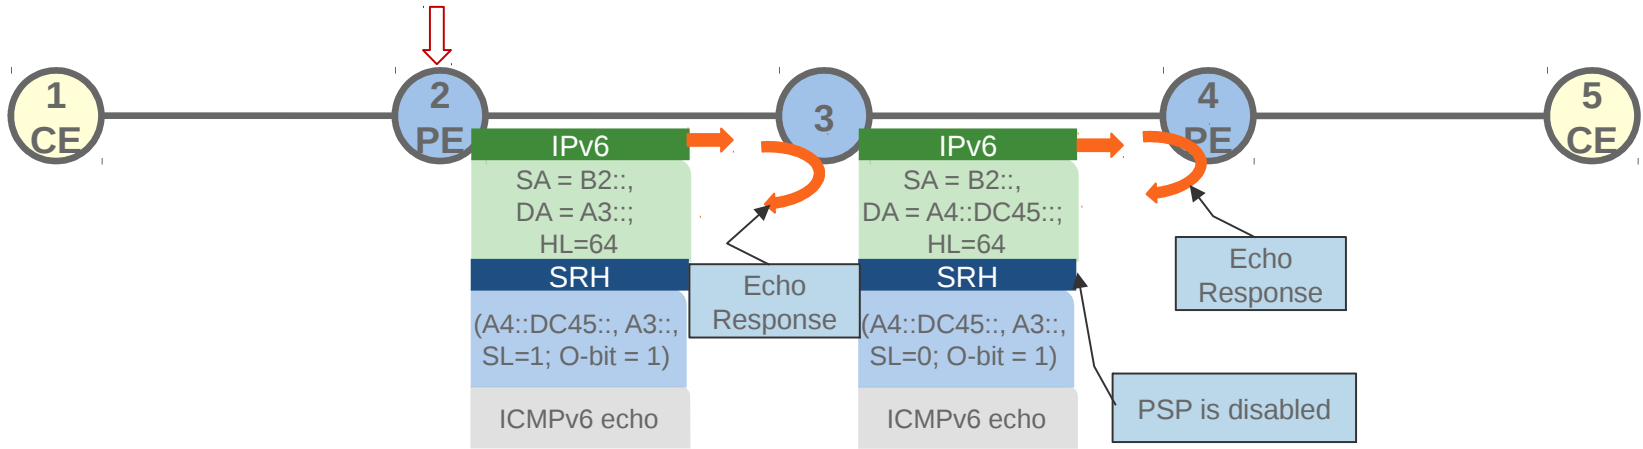
- An:OTP:: SID is instructed in front of the target SID where punt behavior needs to be programmed. E.g., A4::C45 in this example.

Tracing a SID Function (hop-by-hop)



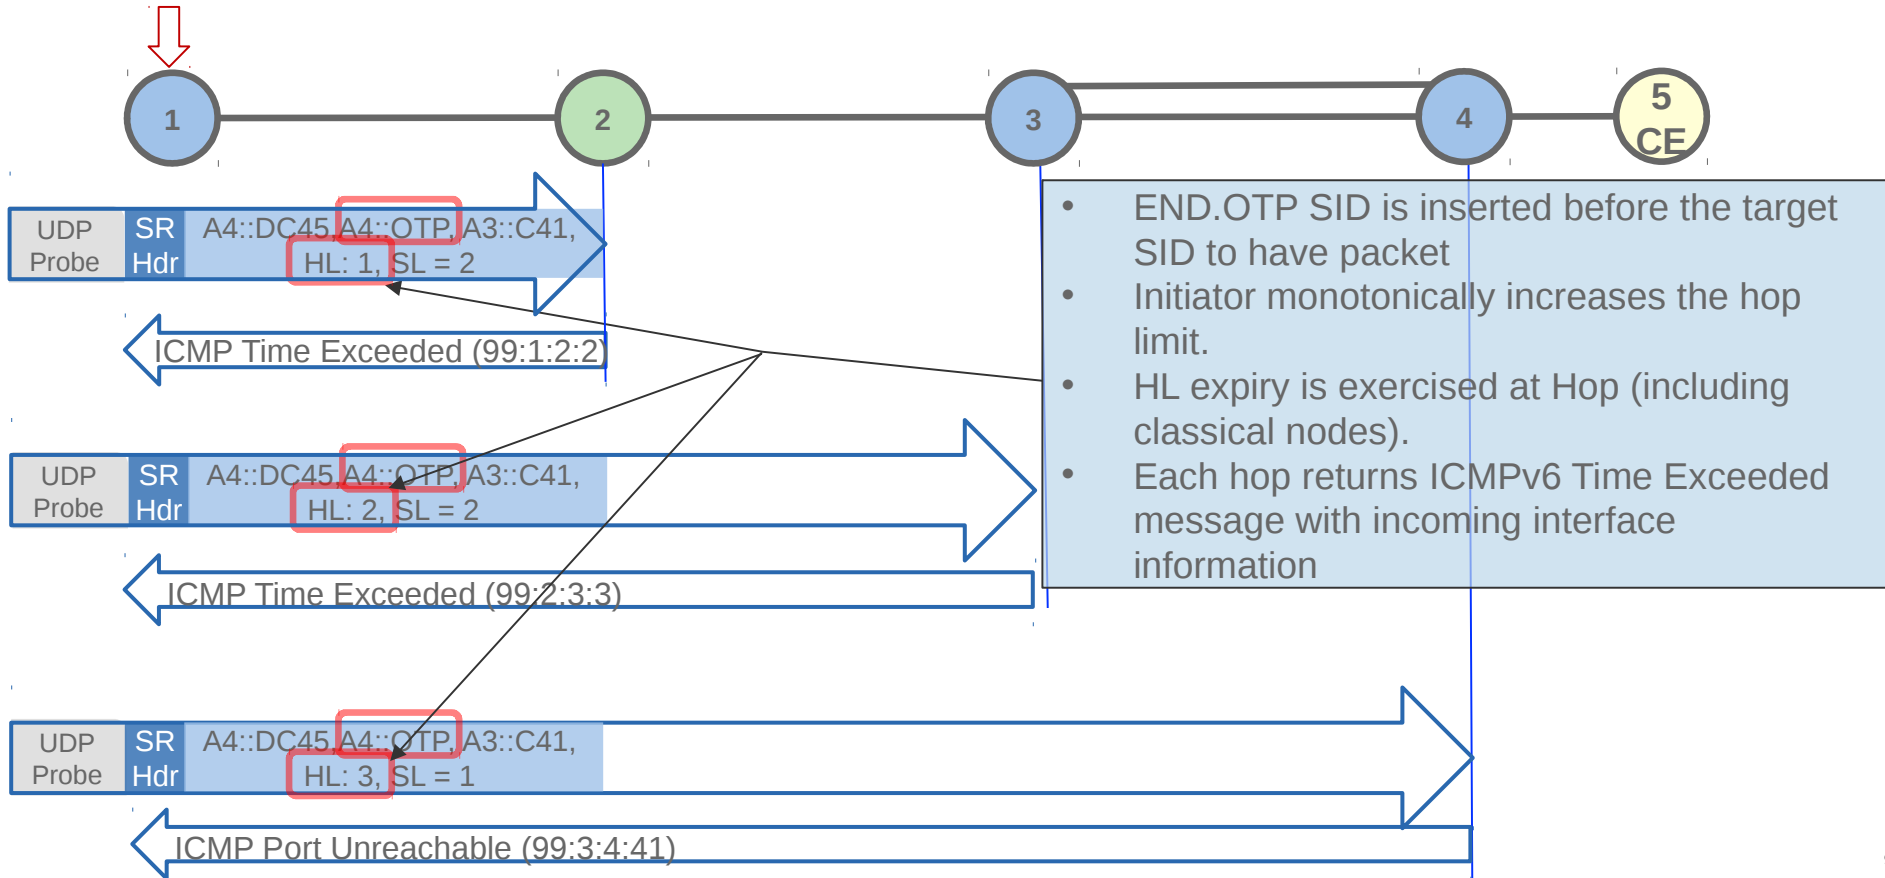
- An::OTP:: SID is instructed in front of the target SID where punt behavior needs to be programmed. E.g., A4::C45 in this example.
 - Due to addition of the An::OTP:: OAM SID to SRH, no change to PSP behavior is required.

Proof of Transit




- Service provider wants to validate if a policy is visiting all segments in the sid-list.
- The ingress injects an out-of-band ping probe with O-bit set in the SRH.

Hop-by-Hop Traceroute




SRv6 Overlay Traceroute



As Hop Limit is set to 64, all classic transit and SRv6 pure transit nodes are skipped in the overlay traceroute.

- O-bit is set and hop limit is set to 64.
- As Hop Limit is set to 64, the classic and SRv6 transit nodes does not respond.
- At each segment node, SRH.Flags.O=1 causes a time-stamped copy of the packet punted and processed.
- Each segment node returns ICMPv6 message with incoming interface information



Next Steps

- Solicit WG review and comments/inputs/feedback.