

Technical Disclosure Commons

Defensive Publications Series

June 2020

SRV6 SID/USID OAM FOR ADJ-SID VALIDATION

Nagendra Kumar Nainar

Carlos M. Pignataro

Zafar Ali

Follow this and additional works at: https://www.tdcommons.org/dpubs_series

Recommended Citation

Nainar, Nagendra Kumar; Pignataro, Carlos M.; and Ali, Zafar, "SRV6 SID/USID OAM FOR ADJ-SID VALIDATION", Technical Disclosure Commons, (June 10, 2020)

https://www.tdcommons.org/dpubs_series/3312



This work is licensed under a [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/).

This Article is brought to you for free and open access by Technical Disclosure Commons. It has been accepted for inclusion in Defensive Publications Series by an authorized administrator of Technical Disclosure Commons.

SRV6 SID/USID OAM FOR ADJ-SID VALIDATION

AUTHORS:

Nagendra Kumar Nainar
Carlos M. Pignataro
Zafar Ali

ABSTRACT

Techniques are presented for using a Loopback segment identifier (SID) that has the semantic of a Loopback address (terminate and process). Such a Loopback SID is assigned for both traditional SID and micro SID (uSID) that can be used in the segment-list of the probe packet to terminate the probe directly in the Nexthop of the adjacent SID (Adj-SID) which can be used to validate the forwarding semantic of the respective Adj-SID at the dataplane.

DETAILED DESCRIPTION

Applying the traditional segment routing over Internet Protocol version 6 dataplane (SRv6) Operations, Administration and Management (OAM) to a node that is configured with an adjacency segment routing identifier (Adj-SID) may not be sufficient to validate the semantic of the Adj-SID configured node. A probe packet can be sent with an O-Flag set in the segment routing header (SRH) or using an OAM-SID. In either the case, the semantic is to punt the probe at the control plane layer, causing the SID to be validated at the control plane and not at the data plane.

Figure 1 below illustrates an example network with node R2 configured/assigned with the following Adj-SID:

- C1::23 towards R3,
- C1::24 towards R4,
- C1::234 --> load balance between R3 and R4.

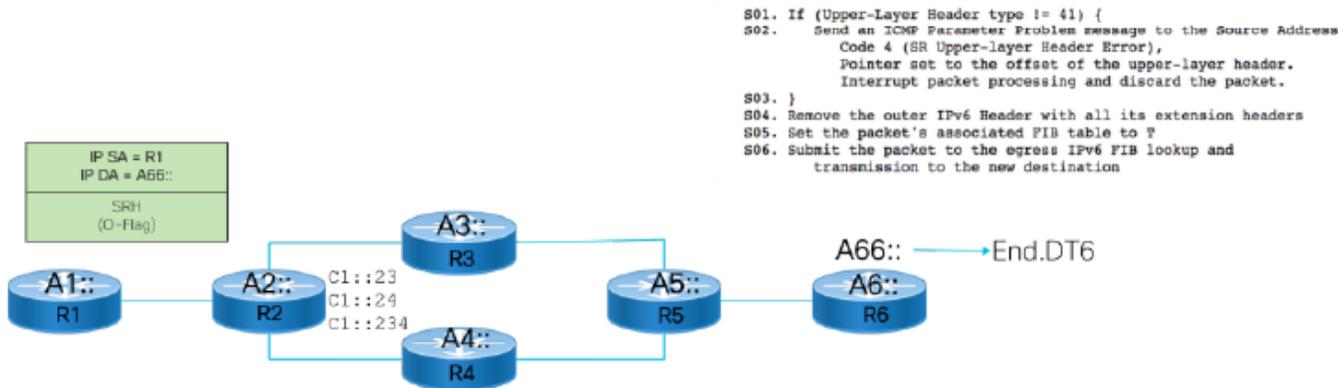


Figure 1

In Figure 1, if node R1 (or any other node) sends a probe for C1::23 with O-Flag or OAM-SID, the probe will end up being punted to a CPU on node R2, resulting in control plane validation but not data plane validation.

When the SID to be validated is the last SID, node R2 might check the NextHeader (which is ICMPv6 in this case) and this may result in a generic parameter problem. One option is to include the prefix-SID of the downstream node (node R3 in this case), however, this may result in a false positive. Thus, including the prefix-SID of the downstream node may be insufficient for validation. For example, if node R2 in the topology of Figure 1 has a forwarding corruption, this will cause any packet with a top SID of C1::23 to POP. The packet may be forwarded to node R4 instead of node R3.

Figure 2 below illustrates a network in which node R1 sends a probe with a segment list (SL) {A3::, C1::23, A2::} to node R2. Node R2 POPs/Rewrites and (wrongly) forwards the probe to node R4. However, node R2 sends the probe with A3::, and the probe may still reach node R3 for validation, resulting in a false positive. While the above problem is described in reference to traditional SID, the same issue is applicable for uSID as well.

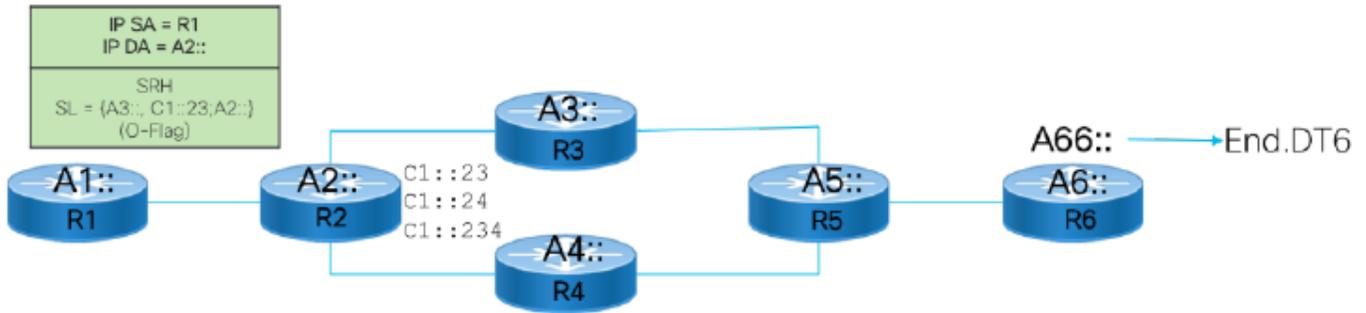


Figure 2

To overcome the above noted deficiencies, a Loopback-SID with a semantic similar to a Loopback address, such as 127.0.0.x (for IPv4) or ::1 (for IPv6), is proposed. In the case of traditional SID scenarios, ::1 may be reused as the Loopback SID. The semantic for traditional SRv6 SID may be as follows:

```
If incoming_sid == ::1 {
  Punt and Process;
}
```

In the case of uSID scenarios, a value can be reserved or configured by a network operator for a Loopback address. For example, a Loopback uSID may be a 16-bit value hex(0001). The destination address may be set to ::1 and thus, the solution works even when the Nexthop (to which the Adj-SID is assigned) is a non-SRv6 node. The semantic for uSID may be as follows:

```
If incoming_sid == hex(0001) {
  //uSID size is 16 bits //
  Punt and Process;
}
```

Figure 3 below illustrates validation behavior with a traditional SID. During validation with a traditional SID, node R1 sends the probe with a SL, in which the last entry is the Loopback SID. Upon receiving the packet at node R2, the top/active SID may be C1::23 (followed by the Loopback SID). Node R2 rewrites the IPv6 destination address with the Loopback SID and forwards the same to the Nexthop node. Any of the Nexthop nodes, upon receiving the probe with destination address as ::1, can punt and process the payload and reply back. When node R2 forwards the packet to the right node (e.g., node R3), node R1 will get a positive response. If node R2 forwards the probe to a wrong node (e.g., node R4), node R1 will get a response from node R4. The response from node R4 can be used to detect the failure.

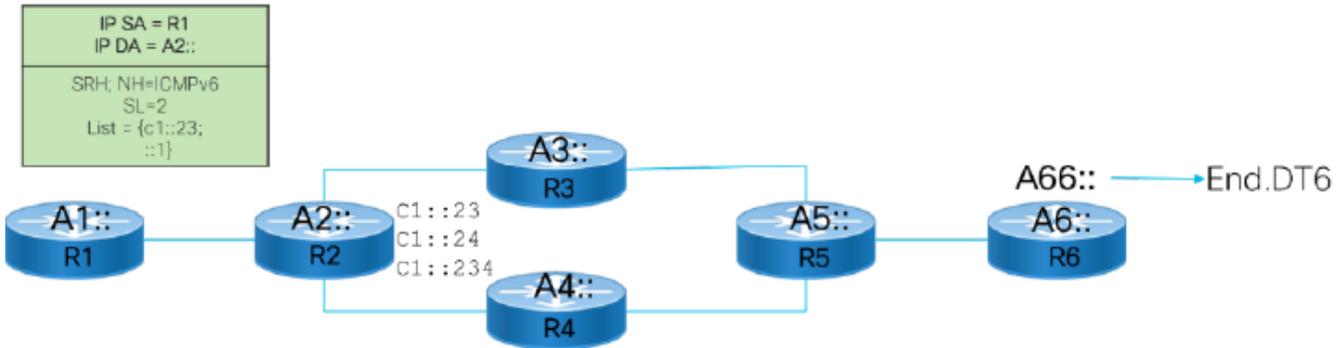


Figure 3

Figure 4 below illustrates validation behavior with a uSID. During validation with a uSID, node R1 sends the probe with <uSID-Block><A2><C23><0001>End-of-Carrier> where C23 is the SID to be validated. Upon node R2 receiving the probe with active SID as C23, node R2 can rewrite the active-SID part with the next SID (which is 0001) and forward to the Nextthop node. The Nextthop node can process the packet and reply back to node R1. The response can be used to detect any mistakes in forwarding.

When the Nextthop is a non-SRv6 node, a tweak of configuring <uSID-Block><Loopback-SID> as one of the loopback addresses on all such non-SRv6 node(s) can be used to terminate the probe.

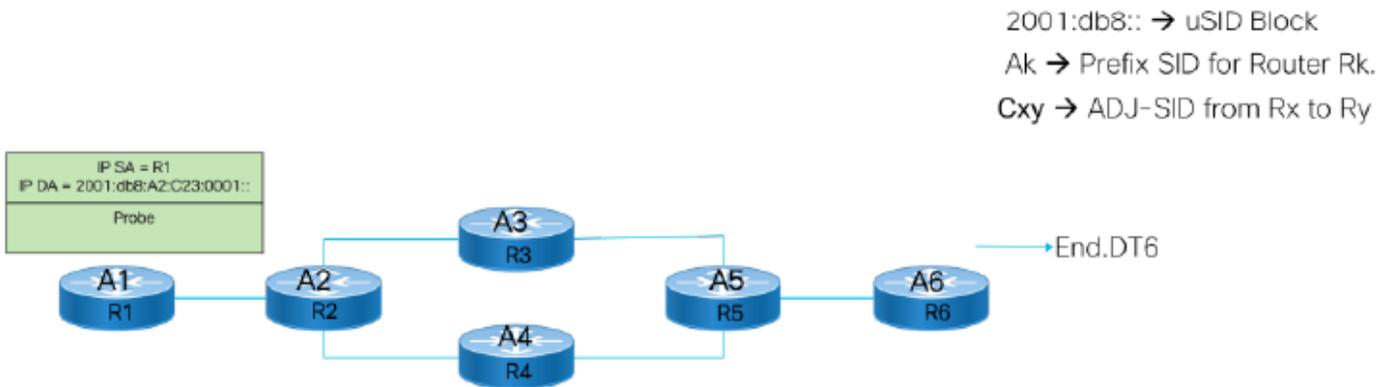


Figure 4

In summary, techniques are described herein for using a Loopback SID that has the semantic of a Loopback address (terminate and process). Such a Loopback SID is assigned for both a traditional SID and a uSID for use in the segment list of the probe packet to

terminate the probe directly in the Nexthop of the Adj-SID. This can be used to validate the forwarding semantic of the respective Adj-SID in the dataplane.