Why are existing mechanisms not enough?

Gyan Mishra

5 mins + 5 mins outrage (65/120)

Why are existing mechanisms not enough? I

1. **DSCP** in the IPv4 and IPv6 Headers [RFC2474]
   - The field is not big enough and only used for Per Hop Behavior QoS scheduling.
   - COS (3bits/7 levels), TOS (3bits/7 levels, DSCP (6bits/64 levels))
2. **IPv6 Flow Label [RFC6437] /MPLS Entropy Label [RFC6790] /Pseudowire Flow Label Stack Entry [RFC6391]**
   - The IPv6 flow label is mainly used for Equal Cost Multipath (ECMP) routing and Link Aggregation [RFC6438].
   - The MPLS entropy label brings a hashable value further up the MPLS label stack
   - [RFC6391] adds a Label Stack Entry (LSE) to facilitate load balancing of the flows within a pseudowire (PW) over the available ECMPs.
3. **SFC ServiceID** [I-D.ietf-sfc-serviceid-header]
   - Subscriber Identifier and Performance Policy Identifier are carried in the Network Service Header (NSH) [RFC8300] Context Header.
   - This is intended only to be used in service function chaining overlays, and carries information between service function nodes.
4. **IOAM Flow ID** [I-D.ietf-ippm-ioam-direct-export]
   - The IOAM Flow ID is used to correlate the exported data of the same flow from multiple nodes and from multiple packets.
   - It is used only within the IOAM structure added to data packets for OAM purposes

Why are existing mechanisms not enough? II

5. **Binding SID** [RFC8402]
   - A BSID is bound to a Segment Routing (SR) Policy and instructs network nodes how to treat a packet
   - BSIDs can only be used in SR networks (SR-MPLS or SRv6)
6. **FlowSpec Label** [RFC5575], [I-D.ietf-idr-flowspec-mpls-match], [I-D.ietf-idr-bgp-flowspec-label], [I-D.liang-idr-bgp-flowspec-route]
   - In BGP VPN/MPLS networks, BGP FlowSpec can be extended to identify and change (push/swap/pop) the labels for traffic that matches a particular FlowSpec rule.
   - Only applies in MPLS networks where BGP is used to distribute the FlowSpec rule bound with labels.
7. **Group Policy ID**
   - The capabilities of the VXLAN-GPE protocol can be extended by defining next protocol "shim" headers that are used to implement new data plane functions.
   - The Group Policy ID is carried in the Group-Based Policy (GBP) Shim header [I-D.lemon-vxlan-lisp-gpe-gbp].
   - GENEVE has similar abilities toVXLAN-GPE to carry metadata.

Gap Analysis

- The existing solutions were all developed for very specific scenarios
  - They have precise and limited functionality
- Each applies to a particular data plane
  - They are not generic across multiple encapsulations and forwarding technologies

- APN aims to define an attribute that:
  - Is generic
  - Can be used for various policy enforcement functions
  - Enables service provisioning
  - Can be carried in all IETF data plane encapsulations

Thank you!