

BGP Update for 5G Edge Service Metadata

draft-dunbar-idr-5g-edge-service-metadata-02

Linda Dunbar: ldunbar@futurewei.com

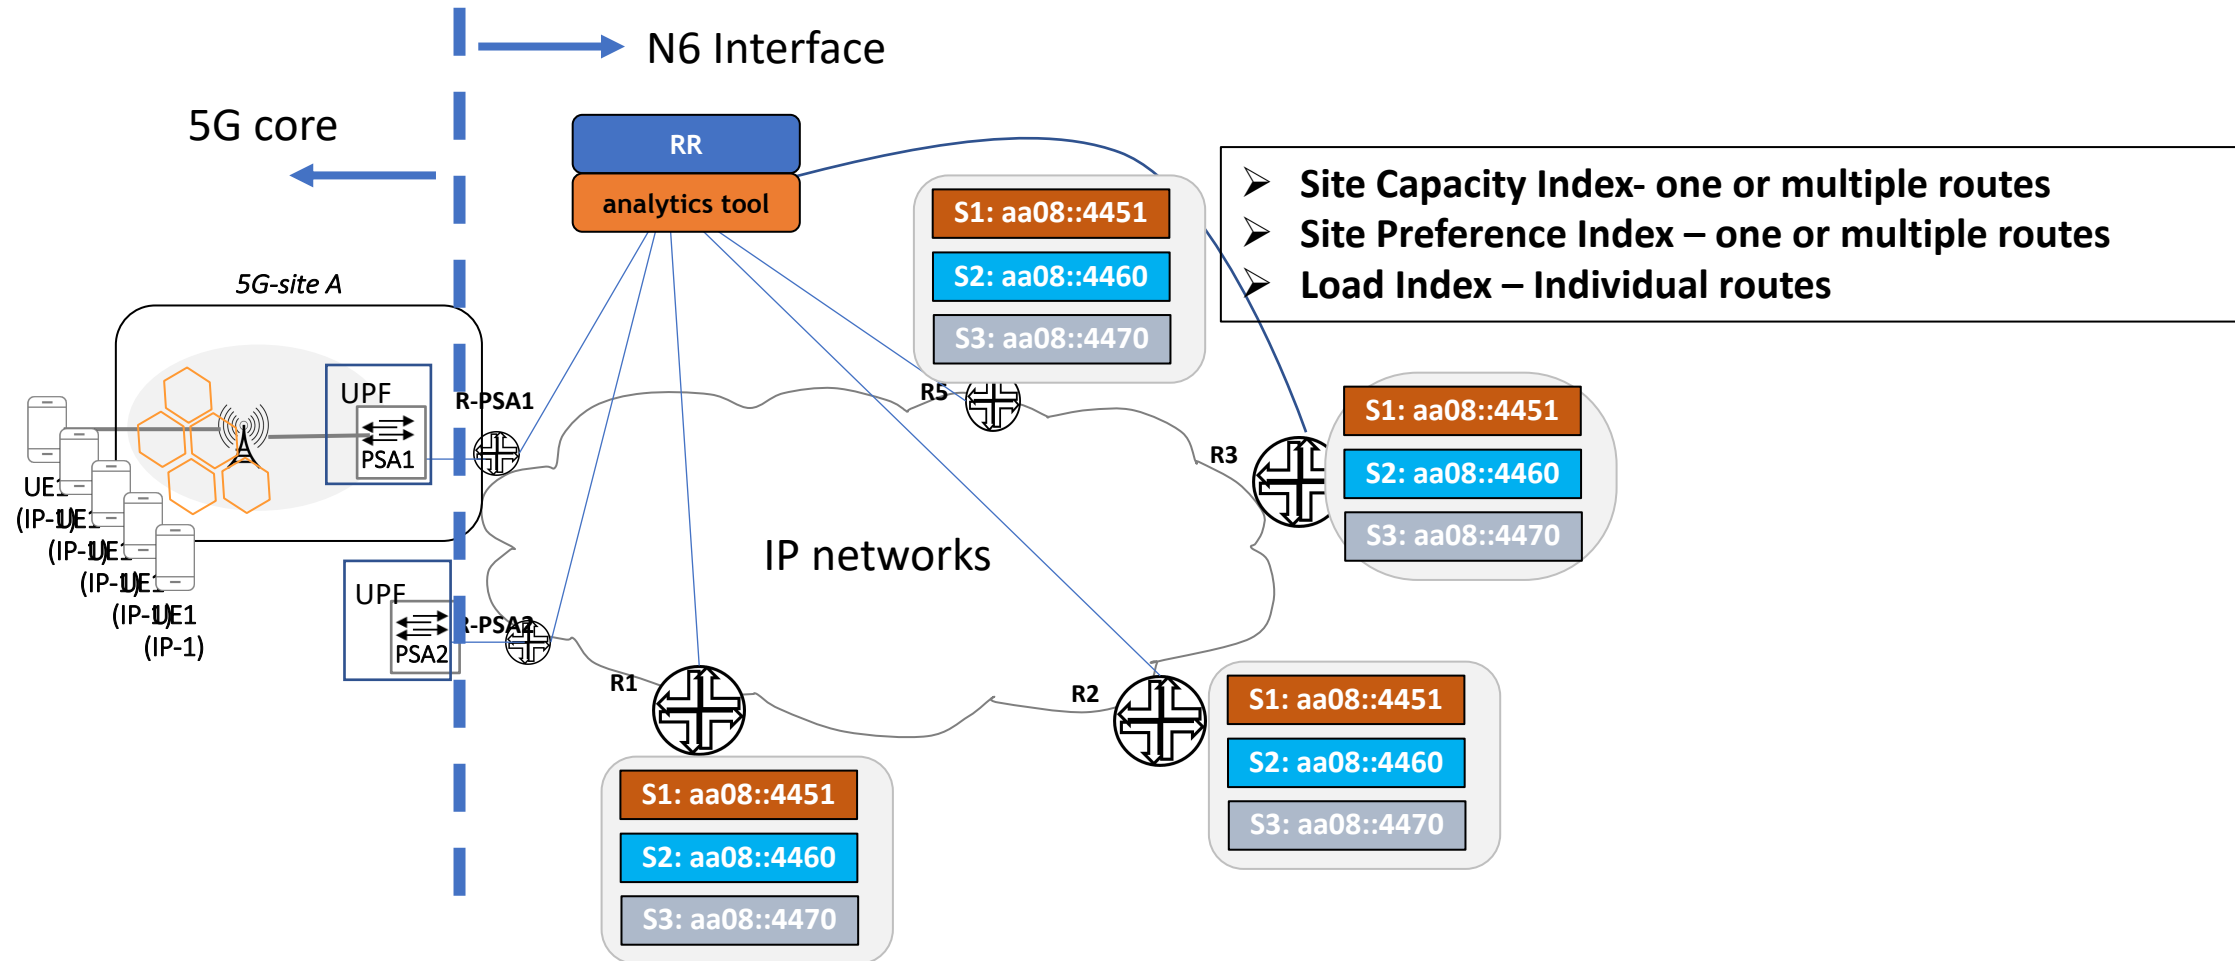
Kausik Majumdar: kmajumdar@microsoft.com

Haibo Wang: Rainsword.wang@huawei.com

Gyan Mishra: gyan.s.mishra@verizon.com


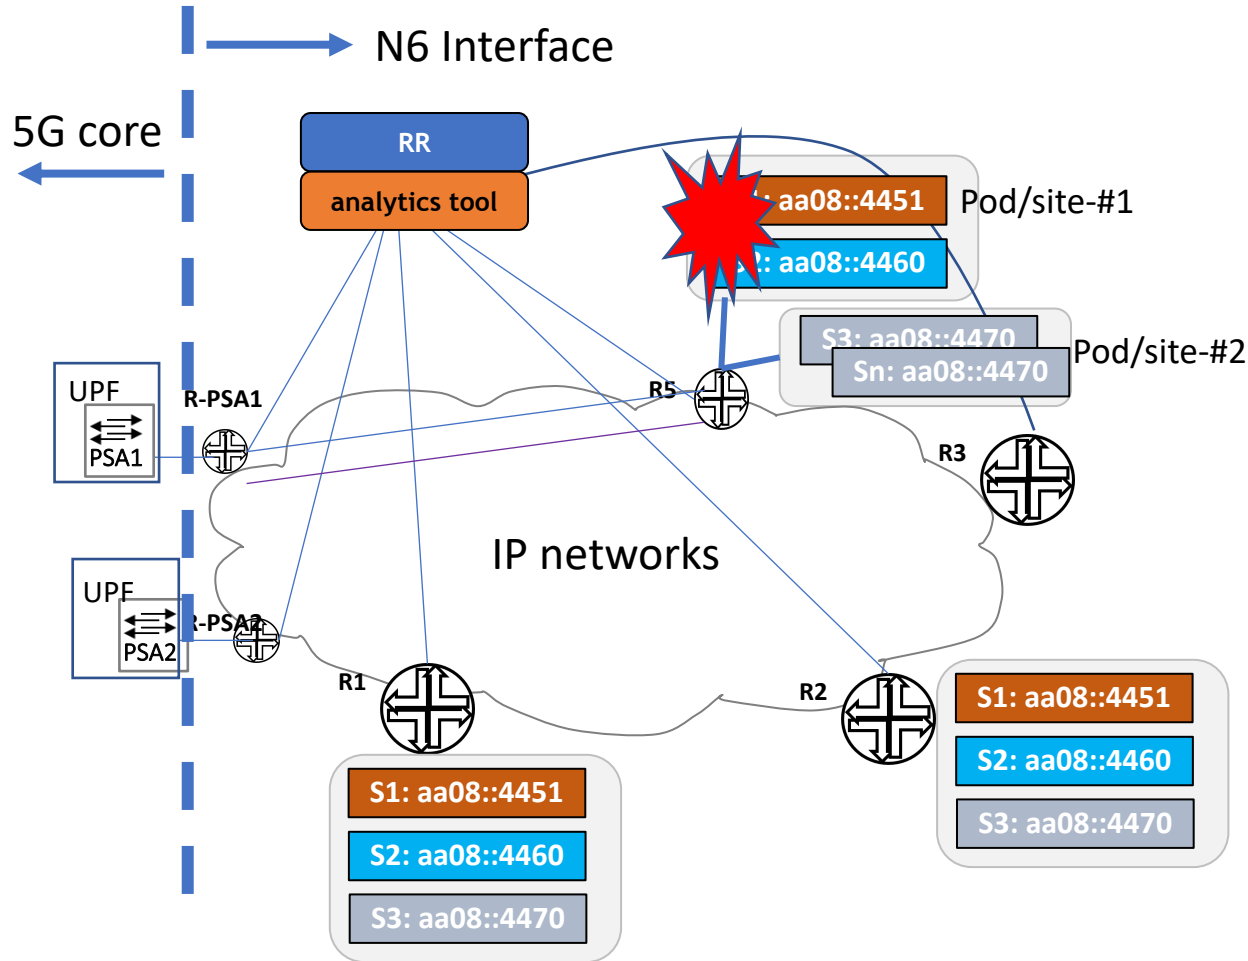
Nov, 2022

Edge Service Metadata: Capacity Index, Site Preference, and Load Index



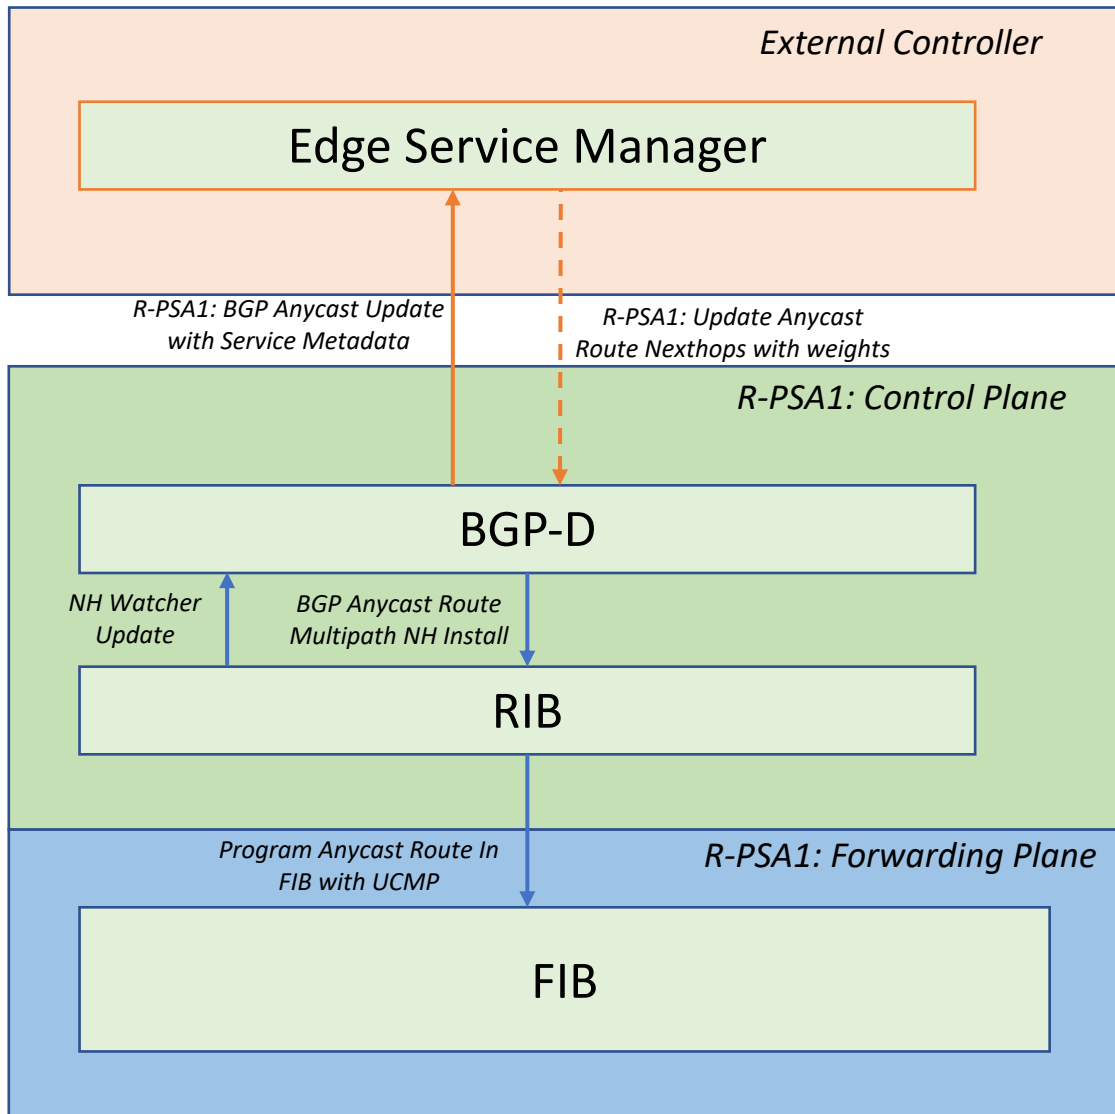
BGP UPDATE for the Site Capacity Index

Capacity Index indicates the capacity value for the site that host the edge services



- Site ID: identifier for a group of routes whose capacity is indicated by the capacity value carried in the UPDATE. There could be more than one sites (or Pods) connected to the egress router (a.k.a. Edge DC GW)
- Site Capacity Index: represent the percentage of the site availability, e.g., 100%, 50%, or 0%. When a site goes dark, the Index is set to 0. 50 means 50% capacity functioning.

Interaction Between BGP and Controller Edge Service Manager



NLRI BGP UPDATE:

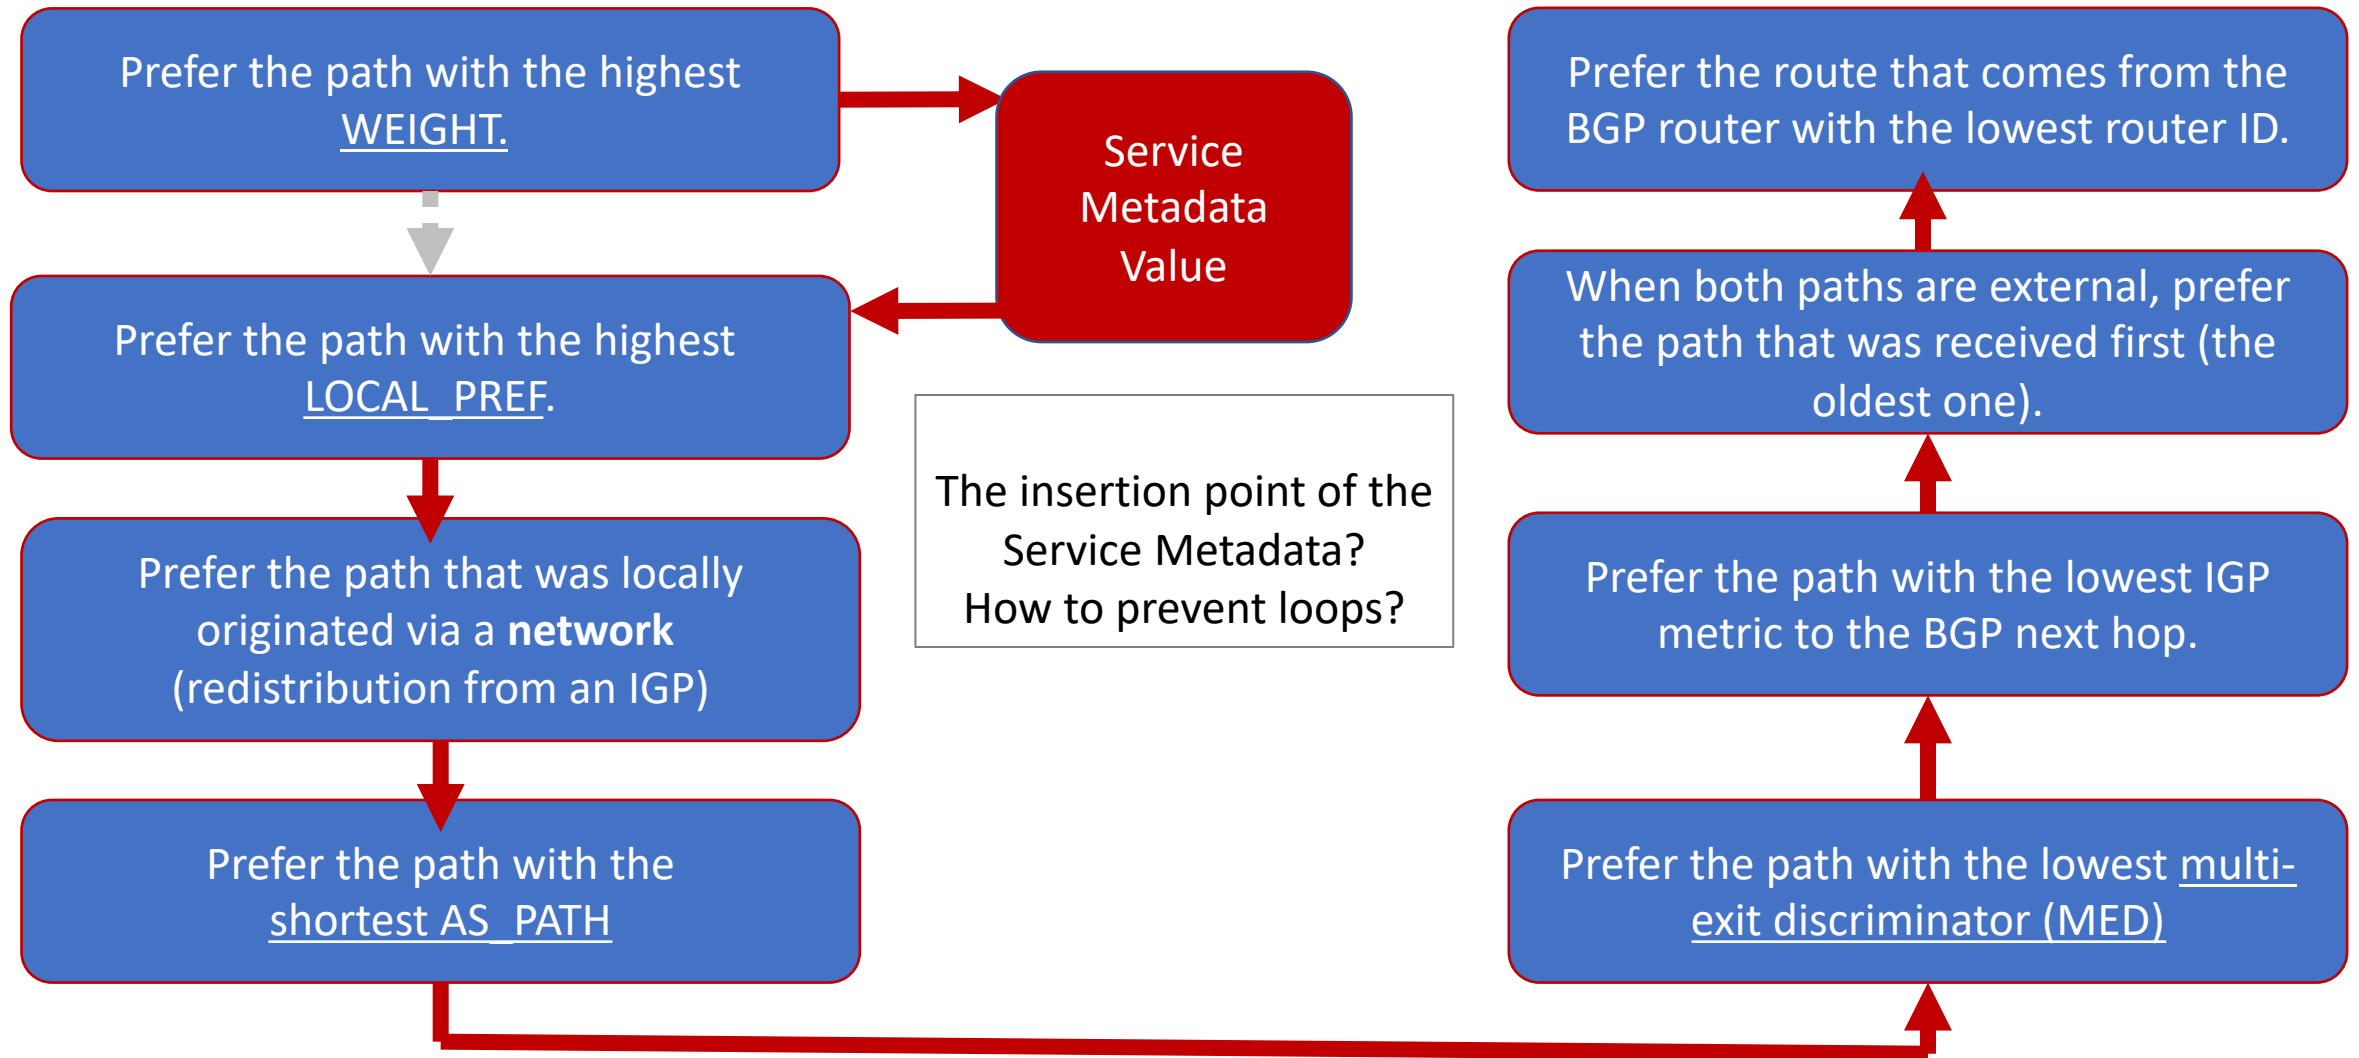
Client route= S1: aa08::4450

Nexthop = [R1, R2, R3, R5]

Service MetaData Attribute:

- Load Measurement sub-TLV
- Capacity Index sub-TLV
- Site Preference sub-TLV

Inserting the Metadata into a chain of factors used in the BGP Best Path Decision process



Unique Properties of the “Next Hop” in 5G Edge Computing network

- Edge Service’s “Next hop” is the exit point (or egress router) of the iBGP domain for the Edge services with the ANYCAST
- The Edge Service Metadata from Egress routers may alter the normal best path decision process
- The iBGP exit points (or egress routers) can be the next hop for many Non-Edge-Service routes.
 - Only a subset of routes with the egress router as Next Hop is considered Edge Services.
 - The Metadata (e.g., Site Capacity Index) should not impact other routes
 - The traditional Weight configured for the next hop only impacts the Non-Edge-service routes.

“Next-Hop” for ANYCAST routes:

At least to have

next-hop-transit that are ASBR within an AS and can propagate NLRI received from others

Add_Path is needed when the NextHop-Transit propagates NLRI from different sources for the same routes

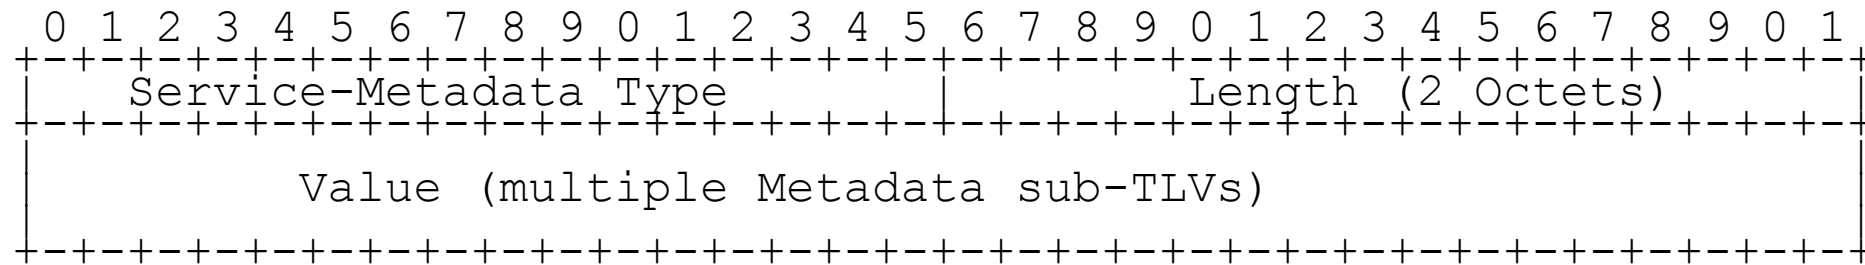
iBGP next-hop-exit-point that has services directly attached.

Add_Path is not needed

Metadata Path Attribute

➤ New Metadata Path Attribute

- optional non-transitive BGP Path attribute to carry the Edge Service Metadata
- The metadata are applicable to multiple NLRIs




➤ Why New Path Attribute, instead of attaching the metadata sub-TLVs to existing path attributes?


- For ease of implementation
- The metadata are applicable to all NLRIs

Edge Service Metadata sub-TLVs


Site Capacity Index: Sub-TLV:



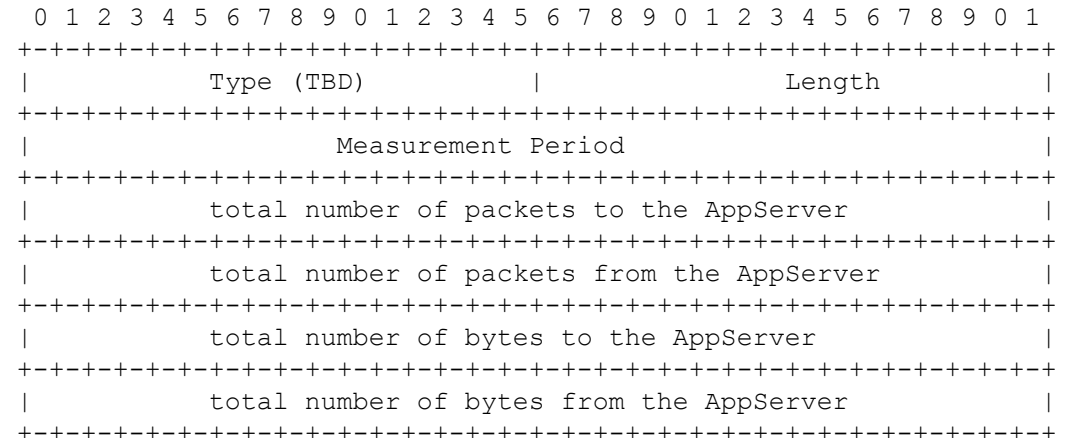
Preference Index sub-TLV:



Aggregated Load Index sub-TLV:



Raw Load Measurement sub-TLV:



Next Step

- Solicit more comments from the WG
- Do we need different definitions:
 - ❖ Next-Hop Transit
 - ❖ Next-Hop iBGP exit points