

BGP Update for 5G Edge Service Metadata

draft-ietf-idr-5g-edge-service-metadata-12

Linda Dunbar: ldunbar@futurewei.com

Kausik Majumdar: kmajumdar@microsoft.com

Haibo Wang: Rainsword.wang@huawei.com

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

Zongpeng Du: duzongpeng@chinamobile.com

IETF 118 Prague

Major changes to Tom Petch comments

- Added many rules for processing Metadata Path Attribute
 - A BGP speaker SHOULD NOT include more than one Metadata Path Attribute in one BGP Update message
 - Emphasize that the four sub-TLV defined in this I-D that there can only be one of each type in a Metadata path attribute. But might not be true for future sub-TLVs specified under the Metadata Path Attribute.
 - The Metadata Path attribute MAY be attached to BGP IPv4/IPv6 Unicast prefixes, BGP Labeled IPv4/IPv6 prefixes [RFC8277], and IPv4/IPv6 Anycast prefixes [RFC4786].
- Error Handling rules:
 - When multiple sub-TLV is present in a Metadata Path Attribute, they are processed independently.
 - When a Metadata Path attribute contains a Sub-TLV whose type is not recognized by a particular BGP speaker; that BGP speaker MUST interpret the attribute as if that Sub-TLV had not been present.
 - Logging the error locally or to a management system is optional.
 - If the “Transitive bit” is set, the unrecognized Sub-TLV remains in the attribute

Major changes to address other mailing list comments

- Resulting next-hop/next-hops ?
 - The Edge Service Metadata is just another attribute in addition to other attributes described in [IDR-CUSTOM-DECISION]
 - aa08::4450 can be reached by three next hops (R1, R2, R3). local BGP's Decision Process based on the traditional network layer policies and metrics identifies the R1 as the optimal next hop
 - the Edge Service Metadata results in R2 as the optimal next hop for the prefix, the Forwarding Plane will have R2 as the next-hop for the destination
- The Edge Service routes are a small number of paid premium services. They should be separated from other regular routes
- RFC4684 used to form the Interested group.
 - The RFC4684 clearly states that Route Target can be an IP address;
 - might need a new Route Target != VPN
- Yao HuiJuan (China Mobile), Yuan Dong Yu (ZTE),
 - Adding clarification that path selection decision is NOT at the Forwarding phase
 - Adding more examples of measurements that can facilitate Service Delay Predication.

IANA Registry

- A new path attribute from the "BGP Path Attributes" registry. The symbolic name of the attribute is "Metadata".

Value	Description	Reference
TDB	Metadata Path Attribute	[this document]

- **Metadata Path Attribute Sub-Types**
 - Registration Procedure: Expert Review

Sub-Type	Description	Reference
0	reserved	[this document]
1	Site Preference Index	[this document]
2	Site Availability Index	[this document]
3	Service Delay Predication	[this document]
4	Raw Load Measurement	[this document]
5-254	unassigned	[this document]
255	reserved	[this document]

Next Step

- Ask for Early Allocation for Metadata Path Attribute
- Working on the second implementation report
- WGLC