BGP Update for 5G Edge Service Metadata

draft-ietf-idr-5g-edge-service-metadata-16

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 119 Brisbane

Metadata Path Attribute

- An optional Non-Transitive BGP Path attribute to carry metrics and metadata about the edge services attached to the egress router
- only a few prefixes BGP advertisement include the metadata path attribute
  - local configuration dictates which prefix has Metadata Path Attribute attached.

GitHub Issue Tracking:

https://github.com/ietf-wg-idr/draft-ietf-idr-5g-edge-service-metadata-14

Changes to Address Jeff Haas Suggestion

- Ingress routers use RFC4684 to register for the interested prefixes that are configured to carry the Metadata Path Attribute.
  - The RFC4684 clearly states that Route Target can be an IP address;
  - Should create a new Route Target (group/VPN, in addition to the IP addresses) ?
- The Security Consideration addition to ensure boundary nodes not leaking Metadata on accident
  - RR attach NO-ADVERTISE well-known community to the UPDATE.

More Changes to Address Jeff Haas' Comments

- Add more clarification for Physical Availability Index Metadata
  - The Site-ID: a group of routes associated with a common physical characteristic, for example a pod, a row of server racks, a floor, or an entire DC.
  - Purpose: one UPDATE message to indicate a group of routes being impacted by a physical event.

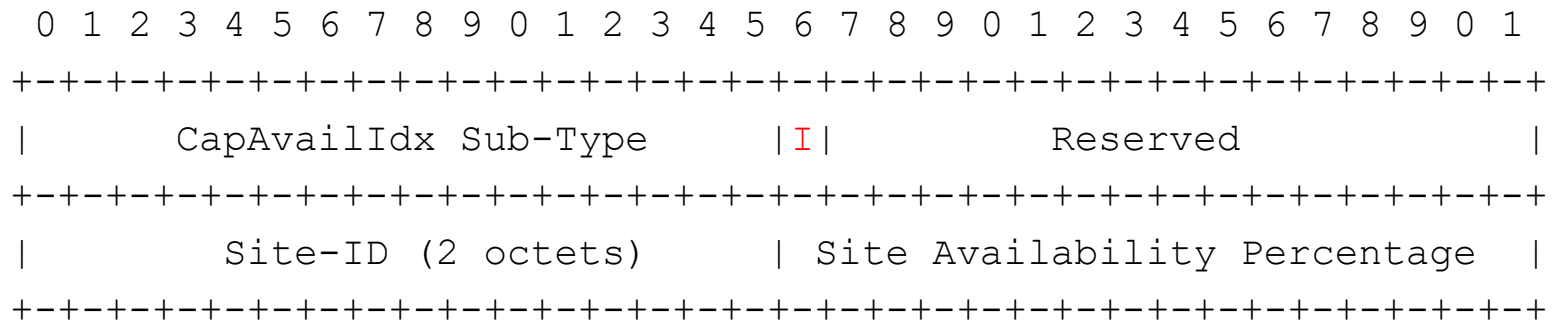

  - RouteFlag I: =1 -> for associating the routes with the Site-ID. The Site Availability Percentage value is ignored.
  - RouteFlag I: =0 -> egress loopback address as NLRI, for receivers to apply the Capacity Availability Index value to all the routes associated with the Site-ID.

Changes to Address Mailing List Comments

- Service Delay caused by Resource Utilization Anomalies.

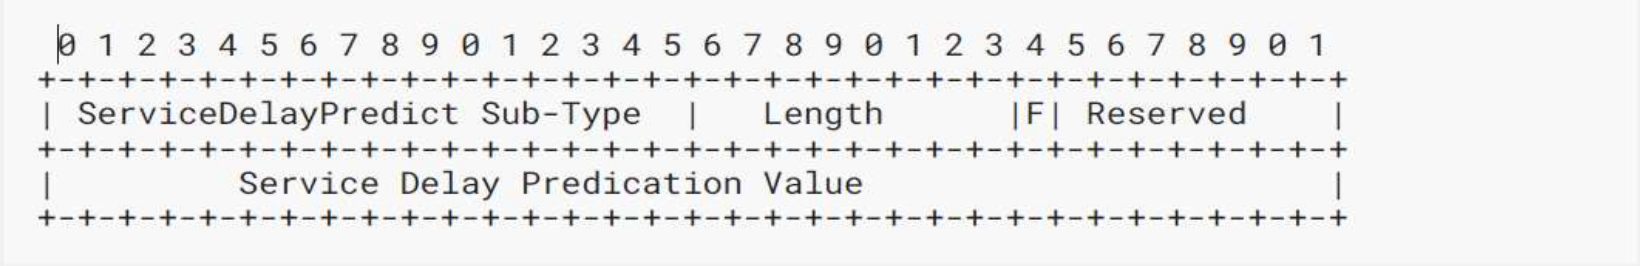

  - Service Delay Predication Value (when the Flag bit is set to 1):an integer in the range of 0-100, with 0 indicating that the service delay is negligible and 100 indicating that the site has the most significant delay compared to all other sites for the same service.
  - Service Delay Predication Value (when the Flag bit is set to 0):the estimated delay time as defined in RFC5905.

Open Issues in Github

- Discussion on previous versions asked about:
  - Route Selection Considerations
  - Route churn
- By going with the standard bgp route selection process (policy + routes), we believe we have avoided this
- Should we delete Section 5 (Service Metadata Influenced Decision Process) as it is implementation details or put in an Appendix?

Next Step

- Close all the open issues in the Github
- Early Allocation
- WGLC

BACKUP SLIDES

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]