Generic Metric extensions for BGP
draft-ietf-idr-bgp-generic-metric-00

09/23 IDR Interim

Srihari Sangli, Juniper Networks
Shraddha Hegde, Juniper Networks
Reshma Das, Juniper Networks
Bruno Decraene, Orange
Bin Wen, Comcast
Mozak Kozak, Comcast
Jie Dong, Huawei
Luay Jalil, Verizon
Ketan Talaulikar, Cisco

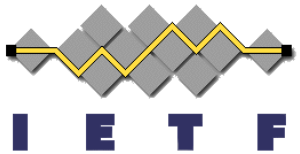

Agenda

- Recap
- Next Hop dependency
- Next steps

Problem statement

- Enables intent-based end-to-end path across multiple AS domains, under common administration
  - IGP-default, Delay, Bandwidth, Other administratively defined metric-types cater various service requirements
- Issues with AIGP (RFC7311)
  - AIGP attribute defined in RFC7311 specifies AIGP TLV to carry default IGP-Metric only
  - Different interpretations of RFC7311 deployed today
  - Discontinuous paths undetected if AIGP were extended to support additional metric types.

Path1: cost=a+b+c+d
Path2: cost=a+e+g

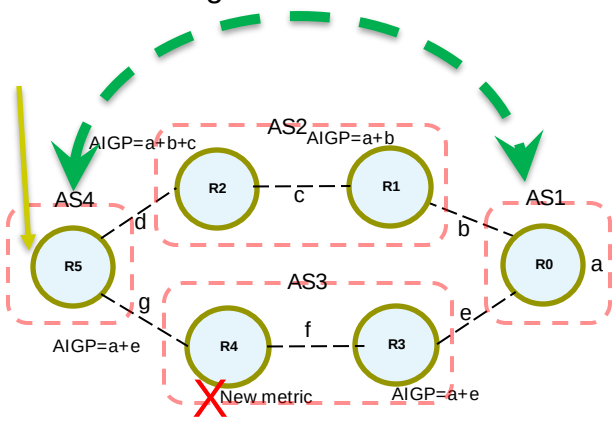

All routers understand new metric except R4

- Network operations can benefit with alignment of metric type & value with IGP registry

Solution recap (1 of 5)

Generic Metric carried in NHC

- Next Hop Capability attribute
  - Optional & transitive
  - Provides next hop level scoping
  - Details in draft-ietf-idr-entropy-label
- Accumulated Metric (AMetric) in NHC
  - Metric type matches with IGP Protocol registry
  - Metric flags supports bits to indicate discontinuity & normalization

Next Hop Capability:

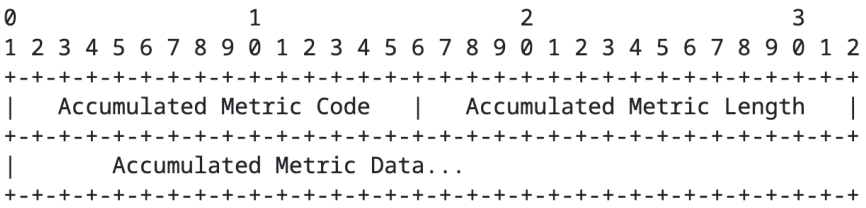

Accumulated Metric Data:

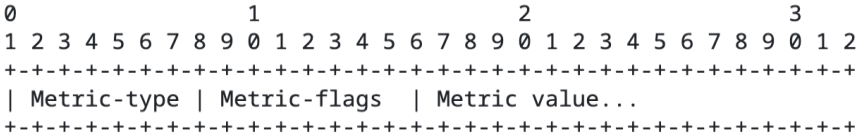

Bit	Description	Usage of Attribute Flags
Bit1: D	Discontinuity	Set to 1 for Discontinuous path
Bit2: N	Normalize	Set to 1 if metric was normalized
Bit 3-8	Reserved	For future use

Solution Recap (2 of 5)

AMetric procedures (draft clarifies and states explicitly)

- If intent is expressed via multiple metric types, the originator router may add more than one AMetric each carrying unique metric.

- For each AMetric, the non-originator router (that modifies next hop) must perform the following
  - Retain received AMetric and update it during reconstruction of NHC
  - Accumulate the metric value with cost to next hop, normalizing the value if metric-type does not match.
  - Regarding metric flags, set “D” to 1 if metric-type is not understood. Set ”N” to 1 if metric normalization was done

Solution Recap (3 of 5)

Handling discontinuity

- Three types of discontinuity
  - Type-A: Router does not support NHC
  - Type-B: Router does not support AMetric
  - Type-C: Router does not support the type of metric representing the end-to-end intent
- NHC attribute and AMetric procedures help detect discontinuity for end-to-end path deterministically
  - Type-A and Type-B handled by NHC procedures
  - Type-C handled by AMetric procedures

Solution Recap (4 of 5)

Best path computation when paths have different metric types

- AMetric (with IGP-default metric) can be compared with AIGP (with IGP-default metric)
- Among paths having AMetric, lower metric-type wins to break the tie
- Implementations should support local policy to specify or override the preferences

Solution Recap (5 of 5)

Deployment considerations

- Discontinuity checks help prevent inadvertent usage of AMetric.
  - AS border routers that support AMetric/base NHC, must filter AMetric at common administration AS boundary
  - If not, the NHC (& AMetric) will leak into neighboring ASes
    - Unless manually filtered
    - or until it reaches an ASBR supporting either AMetric or base NHC
    - NHC detects the leak, filters it and prevents the usage of AMetric by a AMetric capable router
- Ingress router may enforce a policy to handle discontinuous paths
  - Discard / Low preference / Tie-breaker
- All routers in a domain must use AMetric in best path computation unless tunnelling is used to reach the next hop.

Next hop Dependency

- Routers that modify next hop must update Generic Metric information
  - those that do not modify the next hop are only expected to propagate
- Accumulation of Generic Metric across domains is key
  - Provides end-to-end cost for the desired service-intent
- Minimal upgrade in a network is desired
  - Routers that do not modify next hop, especially RR should not require upgrade
- Discontinuity in route accumulation must be detected and acted upon
  - Else it leads to sub-optimal routing

MNH vs NHC

- Next Hop Capability (NHC) attribute is the preferred carrier
  - Built-in Next Hop Scoping
  - Attribute Transitivity and guidance for policy enforcement at the boundary of common administration
  - Simple and extendable to carry information related to next hop
- Multi-Next Hop (MNH) attribute
  - Non-Transitivity brings back limitations of AIGPv1 and hence not preferred
- Recommendation
  - Single or multiple path (via add-path) advertisements for a prefix must use NHC for accumulated Generic Metric
  - If MNH is employed to advertise multiple next hops for a prefix, it must include Generic Metric information in same format as proposed in this document

Next Steps

- Comments welcome

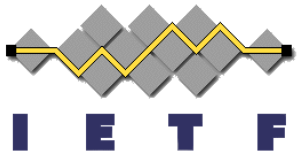

Thank you