

BGP Next-Hop Dependent ~~Capabilities~~ Characteristics Attribute (draft-ietf-idr-entropy-label)

IDR interim meeting, 2024-09-23

John Scudder <jgs@juniper.net>

Brief Reminder (*not* comprehensive)

- NHC attribute is intended to inform downstream routers about (principally) forwarding plane features (“characteristics”).
- Optional transitive, but uses tagging with next hop of originator to ensure the end-to-end path supports the feature in question.
- TLV-based (of course).
 - Current uses: Entropy label (“ELCv3”), Next-next hop nodes (“NNHN”)
 - Pending: Generic metric, IFIT

Changes since last report (so, v10->v15)


- Name changed to “Characteristics” instead of “Capabilities”.
 - Hopefully , this will make it *even clearer* that this is not RFC 5492 Capabilities
- Added section on aggregation rules.
- Many editorial improvements (esp. due to Mach Chen, Ketan Talaulikar reviews).
- Allocated NNHN code point.

Open Issue – link-local

- Per rumor, some implementations support, or might in the future support, peering using *only* link-local addresses (no global address).
- This breaks a correctness assumption, that the BGP next hop can be compared to the next hop in the NHC and it will tell you if the route was altered mid-flight without properly updating the NHC.
- If your implementation doesn't do link-local-only peering, you can ignore the next two slides. 😊

Link-local Example

- Suppose A and D support NHC, B and C don't
- If A sends a route with NHC, and C updates the NH to self before sending to D, D will discard the NHC due to NH mismatch
- But now suppose A peers with B using link-local address L, and C peers with D *also* using link-local address L
- Now D has no way of detecting C set NH to self



Proposed Solution (for v16)

- New TLV that contains (BGP Identifier, ASN) of sender
 - Mandatory if sending a route that has only link-local in its next hop
- When receiver is evaluating a route that has only link-local in its next hop, then it also has to check matching BGP ID, ASN between peer and new TLV
 - If they don't match -- route doesn't match, and NHC is discarded
 - If new TLV is missing -- route doesn't match, and NHC is discarded

Ad-Hoc “Registry”

- Three IDR WG documents, plus one individual submission, use NHCs.
- Allocation policy in IANA section is first come, first served.
- But there’s no IANA registry yet because we haven’t gone to RFC yet.
- Therefore, draft authors are maintaining a “registry” in-line in the draft as part of the IANA section.
- If you have a draft that needs a NHC code point, ask the draft editors to allocate one.

Implementation Status

- It would be *really helpful* if we could move this to publication so we could get IANA to take over registry duties.
- WGLC done.
- Junos implementation done, shipping.
- Need a second one reported.

Thank You!