

BGP Entropy Label Capability, Version 3

draft-scudder-idr-entropy-label-00

IDR, IETF-114, July 27, 2022

John Scudder, Kireeti Kompella, Satya Mohanty, Jim Uttaro, Bin Wen, Serge Krier



Background

- RFC 6790 specifies entropy label, to facilitate load-balancing
 - This is increasingly a must-have for deployments
- It's highly desirable to be able to signal entropy label support
 - The alternative is not using entropy label, or
 - Blindly using entropy label (risks persistent packet loss or misrouting if LSP tail end doesn't support entropy labels)
- RFC 6790 specifies a dataless path attribute that indicates a router can process entropy labels
 - "Entropy Label Capability", or "ELCv1"
 - RFC 6790 requires that the attribute be scoped... but an optional transitive was used

Background [2]

- Juniper developed a solution that fixes the problem, on which this draft is based
 - Documented in draft-scudder-bgp-entropy-label-00
 - We call this “ELCv2”
- Differences:
 - Juniper’s solution reuses Attribute 28
 - The current draft (“ELCv3”) requests a fresh attribute
 - Addresses concerns about ELCv1 and ELCv2 sharing the same attribute codepoint

Requirements

- Constraints for the solution
 - Must not require full route reflector infrastructure upgrade in order to deploy
 - This implies optional transitive
 - Must not leak

Requirements

- Constraints for the solution
 - Must not require full route reflector infrastructure upgrade in order to deploy
 - This implies optional transitive
 - Must not leak
- But these are contradictory. Let's try again:

Requirements

- Constraints for the solution
 - Must not require full route reflector infrastructure upgrade in order to deploy
 - This implies optional transitive
 - Must not leak
- But these are contradictory. Let's try again:
 - ~~Must not leak~~

Requirements

- Constraints for the solution
 - Must not require full route reflector infrastructure upgrade in order to deploy
 - This implies optional transitive
 - Must not leak
- But these are contradictory. Let's try again:
 - ~~Must not leak~~
 - Must do no harm if it does leak

Do no harm

- Approach chosen is to add data to the attribute
 - IP address of the next hop sent in the route by the originator.
- Receiver compares the ELCv3's next hop, to the next hop of the route (the NEXT_HOP or the Network Address of Next Hop field in the MP_REACH_NLRI)
 - If they match, all good
 - If they don't, it was a leak and is discarded

Planned for version 01

- Future extensibility by allowing trailing data
- Considerations for interoperation between ELCv3 and legacy ELCv2
 - Optional of course
 - Essentially, if you receive either a valid ELCv2 or a valid ELCv3, consider the route “EL-capable”

Next steps

- Publish version 01
- Request WG adoption
 - There's a demonstrated need for a standardized solution
 - We have significant deployment experience with ELCv2 which is substantially similar to ELCv3