

# draft-litkowski-bess-vpnv4-ipv6-nh-handling-00

S. Litkowski

S. Agrawal

S. Zhuang

K. Patel

# Problem statement

- Inconsistency in BGP NH encoding brought by RFC5549 compared to RFC4364, RFC4659.

	RFC4364	RFC4659	RFC5549
What is defined ?	VPNv4 over IPv4 network	VPNv6 over IPv4 IPv6 network	VPNv4 over IPv6 network
IPv4 nexthop encoding format	VPN-IPv4	VPN-IPv4	N/A
IPv6 nexthop encoding format	N/A	VPN-IPv6	Regular IPv6 (no zeroes)

- No harm from a pure standardization perspective, BUT...

# Problem statement

- Authors have looked at 9 BGP implementations (including all major vendors) regarding their support of VPNv4 over IPv6 network
- From a BGP running code perspective, we have found:
  - 7 codes following a consistent BGP NH encoding (using VPN-IPv6 address(es))
  - 2 codes not supporting the feature
  - No code compliant with RFC5549

# Well we have an industry problem there...

- There is no reported (AFAWK) interoperability issue
  - The vendors should be interoperable today because they use the same encoding
- There is a standard compliancy issue

# Proposed resolution

- As IETF is driven by running code, let's the standard reflecting running codes
- RFC5549 should be respinned telling that nexthop is encoded using VPN-IPv6 address(es)
- Why ? Moving all the existing running codes to compliancy will just be risk (bugs) and pain (both for vendors and operators)
  - For a zero value !

# Backward compatibility

- Of course, the proposed solution is not backward compatible
- Do we create harm ?
  - Well, today there is no deployment mixing compliant vs consistent implementations
  - There could be a deployment purely compliant with RFC5549 (really unlikely !)
    - We don't break anything here
    - A knob can provide compliancy for future interop requirement with existing consistent implementations
  - There are deployments purely using consistent implementations
    - No change for these deployments

# In addition...

- RFC5549 does not handle the case of SAFI 129
- We would like to add some text to deal with SAFI 129 too (in the same way as SAFI 128)

# Summary

- Our proposal looks reasonable:
  - Accomodating standard to running codes (that's the spirit of IETF)
  - We are happy to hear from any additional existing implementation feedback !
- We need to focus on facts, not just theory !
- We need to create as less pain/risks as we can:
  - Prefer fewer/no code changes industry wide

# Process...

- RFC5549 comes from softwire (INT area)
- VPN SAFI are in BESS
- Authors would like to manage the respin in BESS
- Authors have engaged discussions with ADs. Official AD (INT, RTG) position ?
- In case, we agree to continue the work, what do we do with this document ?
  1. Forget it, and just respin
  2. Progress to RFC as the same time as the respin of RFC5549 for history purpose