

BGP Remote-Next-hop Draft
BGP Tunnel Encapsulation Draft

BGP Remote-Next-hop Draft

- New generic encapsulation attribute
 - Signaled with NLRIs in any address family
 - Attribute carries one or more tunnel endpoints and tunnel parameters
- Helps build dynamic overlay infrastructure
- First presented at IETF, Atlanta 2012
 - WG call requested in July 2014
 - Received a good support for adoption
 - Questions/concerns raised by Eric Rosen and Bruno Decraene

BGP Remote Nexthop Draft

- We have worked through all the concerns
- As a result we now have a new draft: draft-rosen-idr-tunnel-encaps-00
- Moving forward draft-rosen-idr-tunnel-encaps-00 will replace draft-vandevelde-idr-remote-nexthop

Tunnel Encapsulation Draft

- Updates RFC5512 to remove some restrictions
 - Uses Tunnel Encapsulation Attribute defined in RFC5512
 - Tunnel Endpoint MAY not be NLRI of Tunnel SAFI
 - Tunnel Endpoint no longer required to be same IP address as next hop of the update to which it is attached
 - Attribute can now be attached to updates of many different SAFIs, not merely *Encapsulation-SAFI*

What is the Tunnel Encapsulation Attribute?

- Sequence of TLVs, one per tunnel
 - TLV type identifies type of tunnel (e.g., VXLAN, GRE)
 - Sub-TLVs describe how to form encapsulation header when sending packet through the tunnel
 - New *Tunnel Endpoint* sub-TLV defined to allow each tunnels' endpoint address to be explicitly given; endpoint need not be same as next hop of update
 - Per new draft, can be attached to most AFI/SAFIs
 - To forward a packet matching a given route, select one of the tunnels to send it through
 - Choice of tunnel depends on policy

Differences between the *Tunnel Encaps Draft* and the *Remote Next Hop Draft*?

- Technical content of RNH draft largely retained
- Improved integration with RFC5512 mechanisms (less overlap between different mechanisms)
- More detail on the way to use *Encapsulation sub-TLVs* to describe how to form tunnel encapsulation headers
 - Almost enough to permit multi-vendor interop; still more detail needed in some cases
- New sub-TLVs for setting fields of outer IP/UDP headers

More Differences

- New material regarding use of Tunnel Encaps attribute on updates of labeled address families
 - Various options exist for integrating MPLS-based apps with non-MPLS transport tunnels
 - New sub-TLVs are provided allowing one to specify the exact option to be used
- More detail on error handling:
 - When TLVs and/or sub-TLVs are malformed
 - When a TLV is not malformed, but the remote endpoint of the corresponding tunnel is unreachable

More Differences (Cont'd)

- Clarifies that Tunnel Encaps attribute has no effect on BGP bestpath selection
- Addresses issues that can arise when tunnel endpoint is different than update's next hop:
 - Points out some counter-intuitive results that can occur
 - Presents some applicability restrictions, e.g.:
 - No support for “option B” style of VPN interconnect
 - More “don't shoot yourself in the foot” warnings
- Increased emphasis on need for proper scoping of attribute distribution

Next Steps

- Remote-next-hop draft was submitted for WG acceptance
- Some objections were raised
- Tunnel Encaps draft deals with these objections, and is result of collaboration between original authors and some of those who raised objections
- Would like to propose for WG acceptance