BGP Remote Next-Hop

draft-vandevelde-idr-remote-next-hop

G. Van de Velde, K. Patel, D. Roa,
R. Raszuk, R. Bush

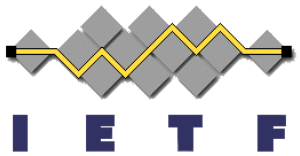

85nd IETF - Atlanta, USA
4 November - 9 November 2012

What is BGP Remote Next Hop?

- New optional transitive attribute
- Facilitate automatic tunneling within and between AS's
- Each attribute carries one or more end-points for an NLRI
- Tunnel encapsulation information is included in the attribute
- (optional RPKI validation can be used for security)

Why this draft even if we have RFC5512

- RFC5512 – Tunnel Encapsulation attribute
- Assumes that tunnel end-point is the BGP speaker
- Is inserted as NLRI in the exchange, while remote-next-hop is an attribute
- Using RFC5512, it is IMPOSSIBLE to set the endpoint to an arbitrary IP address

BGP Remote Next

- One or more BGP next-hops are supported for an exchanged NLRI
- The AF of the exchanged NLRI is independent from the AF of the Remote Next-Hop or set of remote next-hops
- Composed of a set of TLV encodings

TLV Format

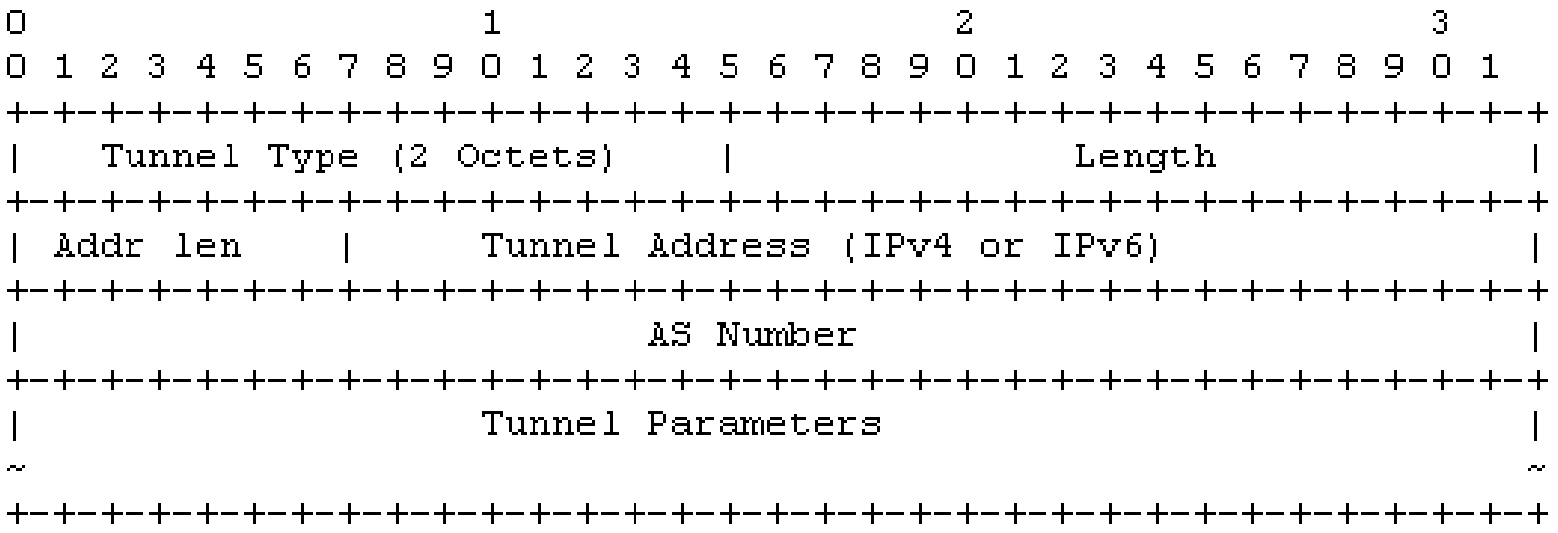

- Tunnel types defined: L2TPv3 over IP, GRE, IP-in-IP, VxLAN, NVGRE
- Tunnel Parameters sub-TLV of RFC5512 is reused

Why?

- Build dynamic overlay Infrastructure
- Multi-homing for IPv6 support
- Virtualization and mobility

Usage note

- Does not need by a capability exchange (hence can be enabled in non-disrupting manner)
- Backwards compatible
- The drafts keeps a place-holder for BGP Next-Hop community to signal valid prefixes as end-points … area that is to be completed

draft-vandevelde-idr-remote-next-hop

THANK YOU!

Backup slides

tbc