IDR
Internet-Draft
Intended status: Standards Track
Expires: September 10, 2015

G. Van de Velde
K. Patel
D. Rao
Cisco Systems
R. Raszuk
NTT MCL Inc.
R. Bush
Internet Initiative Japan

March 9, 2015

BGP Remote-Next-Hop
draft-vandevelde-idr-remote-next-hop-09

Abstract

The BGP Remote-Next-Hop attribute is an optional transitive attribute intended to facilitate automatic tunnelling across an AS for an NLRI in a given address family. The attribute carries one or more tunnel end-points and associated tunnel encapsulation information for a NLRI.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 10, 2015.

Table of Contents

1. Introduction

[RFC5512] defines an attribute attached to an NLRI to signal tunnel end-point encapsulation information between two BGP speakers for a single tunnel. [RFC5512] requires that a new address-family needs to be enabled between the two BGP speakers. It also assumes that the exchanged tunnel endpoint is the NLRI.

This document defines a new BGP transitive attribute known as a Remote-Next-Hop BGP attribute for Intra-AS and Inter-AS usage, and simplifies the exchange and operations involved with tunnel end-point information propagation between two BGP speakers.

The tunnel endpoint information and the tunnel encapsulation information is carried within a Remote-Next-Hop BGP attribute. This attribute can be added to any BGP NLRI. This way the Address Family (AF) of the NLRI exchanged is decoupled from the tunnel SAFI address-family defined in [RFC5512]. Multiple Remote-Next-Hop attribute TLVs can be added to a single NLRI.

Security measures SHOULD be taken to protect against accidental or malicious tampering of the Remote-Next-Hop attribute.

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in [RFC2119] only when they appear in all upper case. They may also appear in lower or mixed case as English words, without any normative meaning.

3. Remote-Next-Hop Attribute

There are an increasing number of use cases where the exchange of multiple unique tunnel endpoints and associated tunnel data is desired for a prefix using segments of an existing infrastructure, where requiring a new address-family to be enabled would add operational complexity.

The BGP Remote-Next-Hop attribute is defined to be attached to each originated BGP NLRI in any applicable address-family. Multiple Remote-Next-Hop attribute TLVs can be applied to a single originated BGP NLRI. Each TLV can contain one or more sub-TLVs that carry encapsulation information. Thus, it enables a simple mechanism to signal multiple, unique tunnel endpoints for a given prefix; as well as multiple encapsulation parameters for prefixes with the same remote tunnel end-point.

BGP Remote-Next-Hop attribute is a Transitive Optional BGP attribute, allowing to signal next-hop encapsulation parameters in a transitive manner without the requirement to enable a new address-family.

This document specifies the tunnel types that can be used with this attribute. The sub-TLVs from [RFC5512] and BGP IPsec tunnel encapsulation [RFC5566] are reused for the BGP Next-Hop-Attribute.

3.1. Tunnel Encapsulation attribute versus BGP Remote-Next-Hop attribute

The use of Tunnel Encapsulation attribute [RFC5512] is based on the principle that the tunnel end-point is carried as part of BGP NLRI in an Encapsulation SAFI.

This requires enabling of the Encapsulation SAFI within a BGP enabled network. It also sets up an interdependency between BGP routes in different SAFIs and the BGP Tunnel SAFI for resolving tunnel next-hops.

The Encapsulation SAFI [RFC5512] assumes that the tunnel endpoint is the NLRI exchanged in the Encaps SAFI, while Remote-Next-Hop decouples the exchanged NLRI from the tunnel end-point information, thereby requiring mutual exclusive usage of the two mechanisms.

While [RFC5512] allows multiple tunnel endpoints and multiple tunnel types to be carried within a BGP Encaps SAFI, the correlation of Tunnel information with other SAFIs is done using the color extended community which is also non-trivial.

4. BGP Remote-Next-Hop attribute TLV Format

This attribute is an optional transitive attribute [RFC1771].

The BGP Remote-Next-Hop attribute is composed of a set of Type-Length-Value (TLV) encodings. The type code of the attribute is (IANA to assign). Each TLV contains information corresponding to a particular tunnel end-point address.

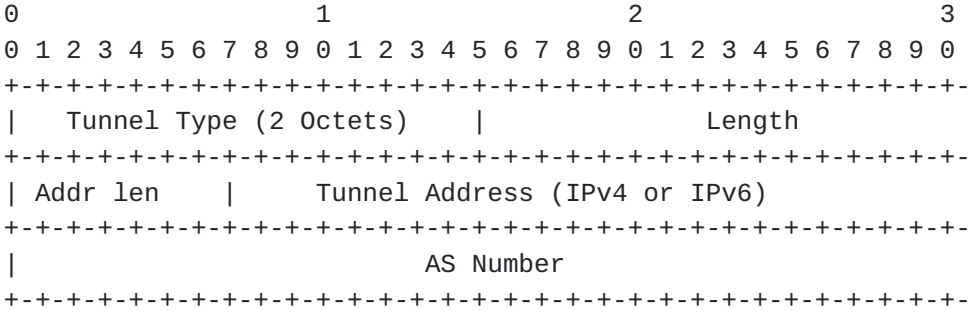

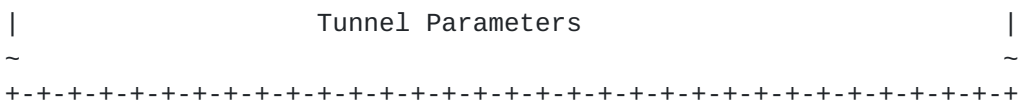

Tunnel Type (2 octets): identifies the type of tunneling technology being signaled. This document specifies the following types:

- L2TPv3 over IP [RFC3931]: Tunnel Type = 1
- GRE [RFC2784]: Tunnel Type = 2
- Transmit tunnel endpoint [RFC5566]: Tunnel Type = 3
- IPsec in Tunnel-mode [RFC5566]: Tunnel Type = 4
- IP in IP tunnel
  with IPsec Transport Mode [RFC5566]: Tunnel Type = 5
- MPLS-in-IP tunnel
  with IPsec Transport Mode [RFC5566]: Tunnel Type = 6
- IP in IP [RFC2003] [RFC4213]: Tunnel Type = 7

This document defines the following types:
- VXLAN: Tunnel Type = 8
- NVGRE: Tunnel Type = 9
- GTP: Tunnel Type = 10
- MPLS-in-GRE: Tunnel Type = 11
- MPLS-in-UDP: Tunnel Type = 12
- MPLS-in-UDP-with-DTLS: Tunnel Type = 13

Unknown types MUST be ignored and skipped upon receipt.

Length (2 octets): the total number of octets of the value field.

Tunnel Address Length (1 octet): Length of Tunnel Address. Set to 4 bytes for an IPv4 address and 16 bytes for an IPv6 address.

AS Number (4 octets): The AS number originating the BGP Remote-Next-Hop attribute and is either a 2-byte AS or 4-Byte AS number

Tunnel Parameter (variable): comprised of multiple sub-TLVs. Each sub-TLV consists of three fields: a 1-octet type, 1-octet length, and zero or more octets of value.

5. Encapsulation sub-TLVs for virtual network overlays

A VN-ID may need to be signaled along with the encapsulation types for DC overlay encapsulations such as [VXLAN] and [NVGRE]. The VN-ID when present in the encapsulation sub-TLV for an overlay encapsulation, MUST be processed by a receiving device if it is capable of understanding it. The details regarding how such a

signaled VN-ID is processed and used is defined in specifications such as [IPVPN-overlay] and [EVPN-overlay].

5.1. Encapsulation sub-TLV for VXLAN

This document defines a new encapsulation sub-TLV format, defined in [RFC5512], for VXLAN tunnels. When the tunnel type is VXLAN, the following is the structure of the value field in the encapsulation sub-TLV:

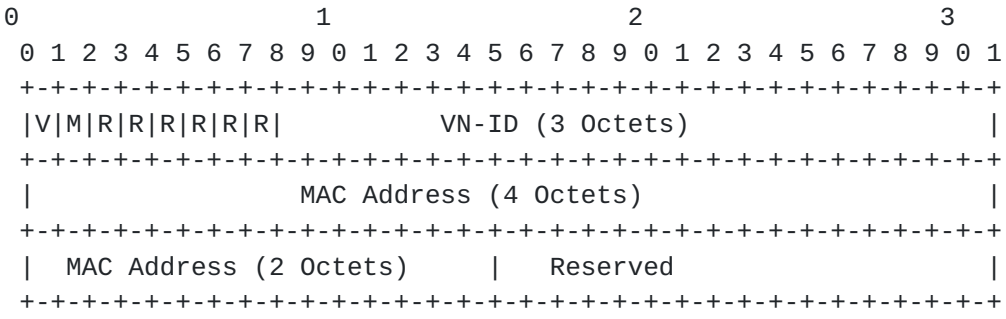

V: When set to 1, it indicates that a valid VN-ID is present in the encapsulation sub-TLV.
M: When set to 1, it indicates that a valid MAC Address is present in the encapsulation sub-TLV.
R: The remaining bits in the 8-bit flags field are reserved for further use. They MUST be set to 0 on transmit and MUST be ignored on receipt.

VN-ID: Contains a 3 octets VN-ID value, if the 'V' flag bit is set. If the 'V' flag is not set, it SHOULD be set to zero and MUST be ignored on receipt.

The VN-ID value is filled in the VNI field in the VXLAN packet header as defined in [VXLAN].

MAC Address: Contains a 6 octets of an Ethernet MAC address if the 'M' flag bit is set. If the 'M' flag is not set, it SHOULD set to all zeroes and MUST be ignored on receipt.

The MAC address is local to the device advertising the route, and should be included as the destination MAC address in the inner Ethernet header immediately following the outer VXLAN header, in the packets destined to the advertiser.

5.2. Encapsulation sub-TLV for NVGRE

This document defines a new encapsulation sub-TLV format, defined in [RFC5512], for NVGRE tunnels. When the tunnel type is NVGRE, the following is the structure of the value field in the encapsulation sub-TLV:

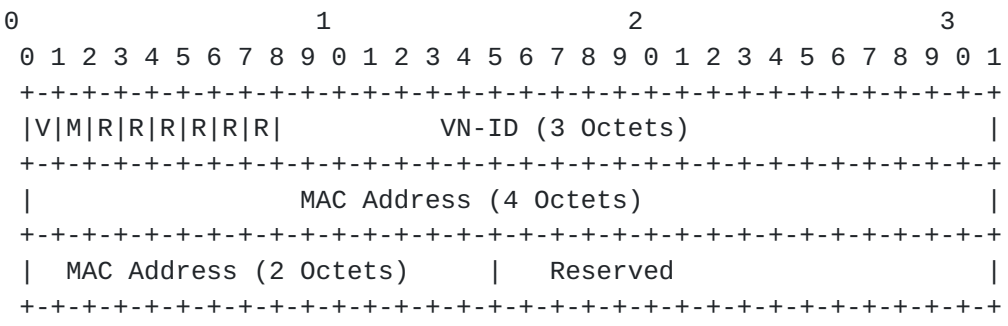

V: When set to 1, it indicates that a valid VN-ID is present in the encapsulation sub-TLV.
M: When set to 1, it indicates that a valid MAC Address is present in the encapsulation sub-TLV.
R: The remaining bits in the 8-bit flags field are reserved for further use. They MUST be set to 0 on transmit and MUST be ignored on receipt.

VN-ID: Contains a 3 octets VN-ID value, if the 'V' flag bit is set. If the 'V' flag is not set, it SHOULD be set to zero and MUST be ignored on receipt.

The VN-ID value is filled in the VSID field in the NVGRE packet header as defined in [NVGRE].

MAC Address: Contains a 6 octets of an Ethernet MAC address if the 'M' flag bit is
set. If the 'M' flag is not set, it SHOULD set to all zeroes and MUST be ignored on receipt.

The MAC address is local to the device advertising the route, and should be included as the destination MAC address in the inner Ethernet header immediately following the outer NVGRE header, in the packets destined to
the advertiser.

5.3. Encapsulation sub-TLV for GTP

This document defines a new encapsulation sub-TLV format, defined in [RFC5512], for GTP tunnels. When the tunnel type is GTP, the following is the structure of the value field in the encapsulation sub-TLV:

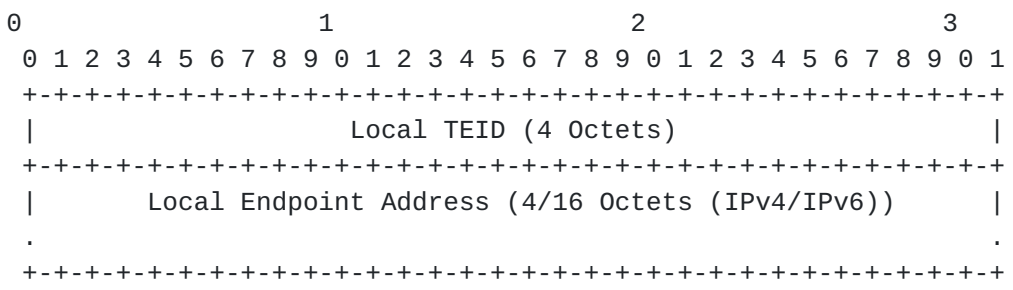

Local TEID: Contains a 32-bit Tunnel Endpoint Identifier of a GTP tunnel assigned by EPC that is used to distinguish different connections in received packets within the tunnel.

Local Endpoint Address: Indicates a 4-octets IPv4 address or 16-octets IPv6 address as a local endpoint address of GTP tunnel.

Local Endpoint Address element makes a tunnel endpoint router allow to have multiple Local TEID spaces. Received GTP packets are identified which tunnel connection by combination of Local Endpoint Address and Local TEID.

5.4. Encapsulation for MPLS-in-GRE

This document defines a new encapsulation sub-TLV format, defined in [RFC5512], for MPLS-in-GRE tunnels. When the tunnel type is MPLS-in-GRE, the following is the structure of the value field in an optional encapsulation sub-TLV:

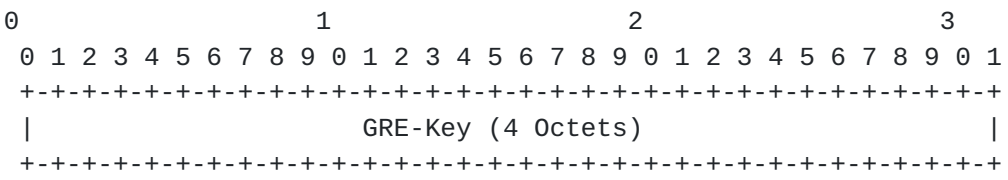

GRE-Key: 4-octet field [RFC2890] that is generated by the advertising router. The actual method by which the key is obtained is beyond the scope of this document. The key is inserted into the GRE encapsulation header of the payload packets sent by ingress routers to the advertising router. It is intended to be used for identifying extra context information about the received payload. Note that the key is optional. Unless a key value is being advertised, the MPLS-in-GRE encapsulation sub-TLV MUST NOT be present.

Note that signaling a GRE tunnel-type with routes in a labeled SAFI may be sufficient to indicate to the receiver that it needs to send MPLS packets with that GRE encapsulation. However, a specific tunnel-type for MPLS-in-GRE is being defined in order to make this indication explicit to a receiver.

6. Remote-Next-Hop Bestpath Considerations

A BGP speaker SHOULD support a policy to enable the support for using BGP Remote Nexthop attribute. An implementation that supports the BGP Remote-Next-Hop MUST use BGP Nexthop attribute information whenever BGP Remote-Next-Hop is not enabled.

Traditionally a BGP speaker uses the IGP cost towards the BGP Next-Hop as a BGP path selection criteria. However, when a BGP speaker is configured to use the BGP Remote-Next-Hop value, then it SHOULD use the IGP cost towards the IP address selected from the Remote-Next-Hop attribute. When there are multiple such IP addresses that may be installed, it SHOULD use the worst IGP cost among them.

Similarly, the speaker SHOULD also check that the IP address is reachable before considering that path eligible for bestpath.

7. Securing Remote-Next-Hop

The Remote-Next-Hop attribute provides a set of tunnel parameters. While the Remote-Next-Hop attribute has as goal to inform an intended recipient with these tunnel parameters, it is important to make sure that the attributes have not been tampered with and that they are restricted to the intended scope of distribution for secure operation.

7.1. Restrictions on Announcing of Remote-Next-Hop Attribute

The Remote-Next-Hop attribute is used to carry an additional information (tunnel end-point, encapsulation type, etc). It has a security value to contain the distribution of the Remote-Next-Hop attribute within its planned scope of distribution. This scope could be, but is not limited to, a particular department, site, organization, across ASes within a same administration control or a global scope.

To contain distribution of the Remote-Next-Hop attribute beyond its intended scope of applicability, attribute filtering MAY be deployed. The BGP speaker communicating to a speaker beyond the intended scope of the Remote-Next-Hop attribute SHOULD filter the attribute during the route announcements.

To facilitate attribute filtering, an implementation that supports the BGP Remote-Next-Hop attribute MUST support a policy to (1) ignore the received attribute and (2) filter the attribute.

7.2. Restrictions on Originating of Remote-Next-Hop Attribute

A BGP Remote-Next-Hop attribute may be added to routes that belong to same Autonomous system as the tunnel endpoint address. Implementations SHOULD validate the following to ensure the validity of Remote-Next-Hop Attribute:

(1) BGP Remote-Next-Hops Tunnel Endpoint and AS number association SHOULD be validated using BGP Origin Validation.

(2) BGP Remote-Next-Hop Tunnel Endpoints underlay routes origin AS SHOULD be validated using BGP Origin Validation. This AS number MUST be the same as the AS number carried within BGP Remote-Next-Hop attribute.

(3) The origin AS of BGP Routes that carry BGP Remote-Next-Hop attribute SHOULD be validated using BGP Origin Validation. This AS number MUST be same as the AS number carried within BGP Remote-Next-Hop attribute.

If the above validation fails, the tunnel type SHOULD be considered as invalid. This does not affect the validity of the others tunnels types carried within the Remote-Next-Hop Attribute.

8. Multiple tunnel endpoint addresses

In some cases, a device may need to accept incoming traffic for a prefix via multiple different encapsulations, to support interactions with remote devices with disjoint capabilities. Certain device implementations cannot support the use of the same IP address as local tunnel endpoint for multiple encapsulations.

In certain cases, a device may need to signal an additional, alternate tunnel endpoint address, to be used by other devices only as a backup in certain failure conditions.

9. Attribute error handling

When receiving a BGP Update message containing a malformed Remote-Next-Hop attribute, the attribute MUST be quietly ignored and not passed along to other BGP peers. (see [draft-ietf-idr-error-handling], Section 7). This is equivalent to the -attribute discard-action specified in [draft-ietf-idr-error-handling]. An implementation MAY log an error for further analysis.

Note that a BGP Remote-Next-Hop attribute MUST NOT be considered to be malformed because it contains more than one TLV of a given type or because it contains TLVs of unknown types.

If a BGP path attribute is received that has the Remote-Next-Hop attribute codepoint but does not have the transitive bit set, the attribute MUST be considered to be a malformed Remote-Next-Hop attribute and MUST be discarded as specified in this section.

10. BGP speakers that do not support BGP Remote-Next-Hop attribute

If a BGP Speaker does not support this attribute, and receives this attribute, then it follows the normal NLRI processing and BGP best path selection, and the resulting forwarding decision is used, as the attribute is optional.

11. Use Case scenarios

This section provides a brief overview of some use-cases for the BGP Remote-Next-Hop attribute. Use of the BGP Remote-Next-Hop is not limited to the examples in this section. Details regarding how the attribute is used are described in the respective solution drafts that are referenced where necessary.

11.1. Stateless user-plane architecture for virtualized EPC (vEPC)

The full usage case of BGP Remote-Next-Hop regarding vEPC can be found in [vEPC], while [RFC6459] documents IPv6 in 3GPP EPS.

3GPP introduces Evolved Packet Core (EPC) that is fully IP based mobile system for LTE and -advanced in their Release-8 specification and beyond. Operators are now deploying EPC for LTE services and encounter rapid LTE traffic growth. There are various activities to offload mobile traffic in 3GPP and IETF such as LIPA, SIPTO and DMM. The concept is similar that traffic of OTT (Over The Top) application is offloaded at entity that is closer to the mobile node (ex. eNodeB or closer anchor).

11.2. Stateless User-plane Architecture for virtual Packet Edge

With the emergence of the NfV technologies, different architectures are proposed for virtualized services. These functions will normally run in the datacenter. BGP remote-next-hop can be used to inject traffic into the virtualized services running in the datacenter using tunnels. These tunnels can be signalled using BGP remote-next-hop. This facilitates a dynamic, simple and clean routing architecture. BGP Remote Next Hop can simplify the orchestration or provisioning layer by signalling the tunnel endpoint (virtual provider edge router) in combination with the encapsulation protocol.

If this is used together with orchestrated traffic steering mechanisms (i.e. BGP Flowspec) , it is possible to differentiate at application level, and forward each different traffic types towards the desired destination.

11.3. Dynamic Network Overlay Infrastructure

The BGP Remote-Next-Hop extension allows consistent signalling of tunnel encapsulations as needed by virtual network overlay solutions such as [I-D.drao-bgp-l3vpn-virtual-network-overlays] and [I-D.sd-l2vpn-evpn-overlay]

11.4. Simple VPN solution using Multi-point Security Association

[draft-yamaya-ipsecme-mpsa] describes the overlay network solution by utilizing dynamically established IPsec multi-point Security Association (SA) without individual connection.

Multi-point SA technology provides the simplified mechanism of the Auto Discovery and Configuration function. This is applicable for any IPsec tunnels such as IPv4 over IPv4, IPv4 over IPv6, IPv6 over IPv4 and IPv6 over IPv6.

MPSA does not provide peer discovery function by itself. However, other mechanism, such as BGP, can be employed with MPSA for automatic peer discovery. BGP Remote-Next-Hop can be used to learn peer information as next-hops.

12. IANA Considerations

This document defines a new BGP attribute known as a BGP Remote-Next-Hop attribute. We request IANA to allocate a new attribute code from the -BGP Path Attributes- registry with a symbolic name -Remote-Next-Hop- attribute.

We also request IANA to allocate four new BGP Tunnel Types from the -BGP Tunnel Encapsulation Attribute Tunnel Types- registry with the following symbolic names: -VXLAN- with Tunnel type 8, -NVGRE- with Tunnel type 9, -GTP- with Tunnel type 10, -MPLS-in-GRE with Tunnel type 11, -MPLS-in-UDP- with Tunnel type 12 and -MPLS-in-UDP-with-DTLS with Tunnel type 13.

13. Security Considerations

This technology could be used as technology as man in the middle attack, however with existing RPKI validation for BGP that risk is reduced.

The distribution of Tunnel end-point address information can result in potential DoS attacks. Therefore is it strongly recommended to install traffic filters, IDSs and IPSs at the perimeter of the tunneled network infrastructure.

measures SHOULD be taken to protect the validity of the BGP Remote-Next-Hop attribute. It is possible to inject a rogue BGP Remote-Next-Hop attribute to an NLRI resulting in Monkey-In-The-Middle attack (MITM). To avoid this type of MITM attack, it is strongly recommended to use a technology mechanism to verify that for NLRI it is the expected BGP Remote-Next-Hop. We anticipate that this can be done with an expansion of RPKI-Based origin validation, see [I-D.ietf-sidr-pfx-validate].

This does not avoid the fact that rogue AS numbers may be inserted or injected into the AS-Path. To achieve protection against that threat BGP Path Validation should be used, see [I-D.ietf-sidr-bgpsec-overview].

14. Privacy Considerations

This proposal may introduce privacy issues, however with BGP security mechanisms in place they should be prevented.

15. Acknowledgements

The authors would like to thanks Satoru Matsushima, Bruno Decraene, Ryuji Wakikawa and Miya Kohno for their usefull vEPC discussions. Istvan Kakonyi provided insight in the vPE use case scenario.

Satoshi Usui provided datapoints around Simple VPN solution using Multi-point Security Association.

16. Change Log

Initial Version: 16 May 2012

Hacked for -01: 17 July 2012

Hacked for -05: 07 January 2014

Hacked for -07: 15 September 2014

17. References

17.1. Normative References

[I-D.ietf-mpls-in-udp]
Xu, X., Sheth, N., Yong, L., Callon, R., and D. Black, "Encapsulating MPLS in UDP", draft-ietf-mpls-in-udp-11 (work in progress), January 2015.

[RFC1771] Rekhter, Y. and T. Li, "A Border Gateway Protocol 4 (BGP-4)", RFC 1771, March 1995.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, March 2000.

[RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.

[RFC3931] Lau, J., Townsley, M., and I. Goyret, "Layer Two Tunneling Protocol - Version 3 (L2TPv3)", RFC 3931, March 2005.

[RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, October 2005.

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

[RFC5512] Mohapatra, P. and E. Rosen, "The BGP Encapsulation Subsequent Address Family Identifier (SAFI) and the BGP Tunnel Encapsulation Attribute", RFC 5512, April 2009.

[RFC5566] Berger, L., White, R., and E. Rosen, "BGP IPsec Tunnel Encapsulation Attribute", RFC 5566, June 2009.

[RFC6459] Korhonen, J., Soininen, J., Patil, B., Savolainen, T., Bajko, G., and K. Iisakkila, "IPv6 in 3rd Generation Partnership Project (3GPP) Evolved Packet System (EPS)", RFC 6459, January 2012.

[RFC7348] Mahalingam, M., Dutt, D., Duda, K., Agarwal, P., Kreeger, L., Sridhar, T., Bursell, M., and C. Wright, "Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", RFC 7348, August 2014.

17.2. Informative References

[I-D.drao-bgp-l3vpn-virtual-network-overlays]
Rao, D., Mullooly, J., and R. Fernando, "Layer-3 virtual network overlays based on BGP Layer-3 VPNs", draft-drao-bgp-l3vpn-virtual-network-overlays-03 (work in progress), July 2014.

[I-D.ietf-idr-error-handling]
Chen, E., Scudder, J., Mohapatra, P., and K. Patel, "Revised Error Handling for BGP UPDATE Messages", draft-ietf-idr-error-handling-13 (work in progress), June 2014.

[I-D.ietf-sidr-bgpsec-overview]
Lepinski, M., "An Overview of BGPsec", draft-ietf-sidr-bgpsec-overview-06 (work in progress), January 2015.

[I-D.ietf-sidr-pfx-validate]
Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", draft-ietf-sidr-pfx-validate-10 (work in progress), October 2012.

[I-D.matsushima-stateless-uplane-vepc]
Matsushima, S. and R. Wakikawa, "Stateless user-plane architecture for virtualized EPC (vEPC)", draft-matsushima-stateless-uplane-vepc-01 (work in progress), July 2013.

[I-D.sd-l2vpn-evpn-overlay]
Sajassi, A., Drake, J., Bitar, N., Isaac, A., Uttaro, J., and W. Henderickx, "A Network Virtualization Overlay Solution using EVPN", draft-sd-l2vpn-evpn-overlay-03 (work in progress), June 2014.

[I-D.sridharan-virtualization-nvgre]
Sridharan, M., Greenberg, A., Wang, Y., Garg, P., Venkataramiah, N., Duda, K., Ganga, I., Lin, G., Pearson, M., Thaler, P., and C. Tumuluri, "NVGRE: Network Virtualization using Generic Routing Encapsulation", draft-sridharan-virtualization-nvgre-05 (work in progress), July 2014.

[I-D.yamaya-ipsecme-mpsa]
Yamaya, A., Ohya, T., Yamagata, T., and S. Matsushima, "Simple VPN solution using Multi-point Security Association", draft-yamaya-ipsecme-mpsa-04 (work in progress), July 2014.

Authors' Addresses

Gunter Van de Velde

Email: gunter@vandevelde.cc

Keyur Patel
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95124 95134
USA

Email: keyupate@cisco.com

Dhananjaya Rao
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95124 95134
USA

Email: dhrao@cisco.com

Robert Raszuk
NTT MCL Inc.
101 S Ellsworth Avenue Suite 350
San Mateo, CA 94401
US

Email: robert@raszuk.net

Randy Bush
Internet Initiative Japan
5147 Crystal Springs
Bainbridge Island, Washington 98110
US

Email: randy@psg.com