Internet Draft Expires: March 2006

Internet Draft Expires: March 2006 Ruchi Kapoor Gargi Nalawade Chandra Appanna Cisco Systems

BGPv4 Tunnel Encapsulation Attribute

draft-kapoor-nalawade-idr-bgp-ssa-02.txt

<u>1</u>. Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with <u>Section 6 of BCP 79</u>.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/lid-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at <u>http://www.ietf.org/shadow.html</u>.

2. Copyright Notice

Copyright (C) The Internet Society (2005). This document is subject to the rights, licenses and restrictions contained in $\frac{BCP}{78}$, and except as set forth therein, the authors retain all their rights."

"This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

3. Abstract

This document defines a new BGP attribute called the Tunnel Encapsulation Attribute. This attribute is used to carry Tunnel encapsulations, capabilities and other relevant information about a Tunnel.

<u>4</u>. Introduction

In case of MPLS-VPNs, the reachability to the nexthop is provided via LDP, which provides information about the LSP Tunnel endpoint. However, if the reachability were to be via a Tunnel, then there is no mechanism to be provide that reachability dynamically. One could manually configure Tunnels on each PE and use policy to specify what Tunnel to use for reaching what BGP nexthop. However this is too complex and unmanageable. What is needed is a way to send the Tunnel Encapsulation information dynamically.

This document defines a new BGP attribute called the Tunnel Encapsulation Attribute which will be used to carry Tunnel Encapsulations with BGP updates. The value part of the Tunnel Encapsulation attribute can contain one or more Tunnel Encapsulation TLVs. The TLVs can carry Tunnel Encapsulations for one or more Tunnel types originating on a given router. The Type in the TLV identifies what kind of Tunnel it is.

Also, unlike the fixed length (8 octets) extended community types, the Tunnel encapsulation TLV is a variable length entity, so it can carry data of any length, as required by the specific type of a Tunnel Encapsulation TLV.

5. Format of the Tunnel Encapsulation Attribute

draft-kapoor-nalawade-idr-bgp-ssa-02.txt
2]

Internet Draft 2005

The format of the BGP Tunnel Encapsulation attribute will be as follows:

The BGP Tunnel Encapsulation Attribute is a variable length, optional transitive attribute. The Value field of the attribute may contain one or more Tunnel Encapsulation TLVs. The Value field of the Tunnel Encapsulation Attribute is encoded as follows and may contain one or more tuples of the following:

Type Field : 2 octetsLength Field: 2 octetsValue Field : Variable

Type: The Type field identifies the type of the Tunnel Eg. mGRE or IPSEC. The format of the Type Field is as shown below:

T - Transitive bit

draft-kapoor-nalawade-idr-bgp-ssa-02.txt
3]

Internet Draft 2005

If the 'T' bit has a value of 1, it implies that the Tunnel type is transitive across ASes. If the 'T' bit has a value of 0, it implies that the Tunnel type is non-transitive across ASes.

Remaining 15 bits: Indicate the type of the TLV.

Length: The Length is 2 octets long and indicates the length of the Value field.

Value: The Value fields for all Tunnel TLVs, will carry the following fields :

1 Preference (2 octets) |S| Flags | +-+-+-+-+-+-+-+ + | Tunnel Encapsulation (variable)

where

Preference - Preference is a 2 Octet field containing a Preference associated with the TLV. The Preference value indicates a preferred ordering of tunneling encapsulations according to the sender (i.e. egress PE). The recipient of the information SHOULD take the sender's preference into account in selecting which encapsulation it will use. A higher value indicates a higher preference.

Flags - Flags is a 1 Octet field containing flag-bits. The leftmost bit indicates whether Sequence numbering is to be used or not. The remaining bits are reserved for future use.

Tunnel Encapsulation: The Tunnel Encapsulation field will carry the tunnel encapsulation information, specific for this particular tunnel 'Type'.

The Value field will have a fixed part and an optional variable part as specified by the respective Tunnel Types. The variable part when present will contain Sub-TLVs encoded as follows :

Sub-Type Field: 1 octetsLength Field : 1 octetsValue Field : Variable

draft-kapoor-nalawade-idr-bgp-ssa-02.txt
4]

- 5 -

6. Operation

A BGP speaker that supports the Tunnel Encapsulation Attribute, MUST accept all Tunnel Encapsulation TLV types. If the Tunnel Encapsulation TLV is transitive, The BGP speaker MUST forward it even if it does not understand it. This attribute MUST only be carried with the BGP Tunnel SAFI [BGP-TUN-SAFI].

If this attribute is carried with another BGP SAFI, the receiving BGP Speaker MUST ignore this attribute.

A Capability has not been defined for this attribute intentionally. The presence of the Tunnel SAFI implies the Capability to understand this Tunnel Encapsulation attribute.

7. Security Considerations

This extension to BGP does not change the underlying security issues.

8. Acknowledgements

We would like to thank Dan Tappan and Jim Guichard for their significant contributions. We would also like to thank Francois LeFaucher, Arjun Sreekantiah, Shyam Suri and John Scudder for their comments and suggestions.

9. Normative References

[BGP-4] Rekhter, Y. and T. Li (editors), "A Border Gateway Protocol 4 (BGP-4)", Internet Draft <u>draft-ietf-idr-bgp4-26.txt</u>, April 2005.

[BGP-CAP] Chandra, R., Scudder, J., "Capabilities Advertisement with BGP-4", <u>draft-ietf-idr-rfc2842bis-02.txt</u>, April 2002.

[IANA-AFI] <u>http://www.iana.org/assignments/address-family-numbers</u>.

[IANA-SAFI] <u>http://www.iana.org/assignments/safi-namespace</u>.

10. Author's Addresses

Ruchi Kapoor mailto:ruchi@cisco.com

Gargi Nalawade mailto:gargi@cisco.com

Chandra Appanna mailto: achandra@cisco.com

draft-kapoor-nalawade-idr-bgp-ssa-02.txt
5]

Cisco Systems, Inc 170 West Tasman Drive San Jose, CA 95134

<u>11</u>. Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in <u>BCP 78</u> and <u>BCP 79</u>.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

<u>12</u>. Full Copyright Statement

Copyright (C) The Internet Society (2005). All Rights Reserved.

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED draft-kapoor-nalawade-idr-bgp-ssa-02.txt
6]

Internet Draft 2005

WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

<u>13</u>. Expiration Date

This memo is filed as <<u>draft-kapoor-nalawade-idr-bgp-ssa-02.txt</u>> expires March, 2006. draft-kapoor-nalawade-idr-bgp-ssa-02.txt
7]