

Networking Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 12, 2007

F. Le Faucheur
Cisco Systems
E. Rosen
Cisco Systems, Inc.
October 9, 2006

**Advertising an IPv4 NLRI with an IPv6 Next Hop
draft-lefaucheur-idr-v4nlri-v6nh-00.txt**

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 [Section 6 of BCP 79](#).

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/1id-abstracts.txt>.

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

This Internet-Draft will expire on April 12, 2007.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

MultiProtocol-BGP (MP-BGP) specifies that the set of Network Layer protocols to which the address carried in the Next Hop field may belong is determined by the Address Family Identifier (AFI) and the Subsequent Address Family Identifier (SAFI). The current AFI/SAFI definitions for the IPv4 address family only have provisions for advertising a Next Hop address that belongs to the IPv4 protocol when advertising an IPv4 Network Layer Reachability Information (NLRI).

This document specifies the extensions necessary to allow advertising an IPv4 NLRI with a Next Hop address that belongs to the IPv6 protocol. This comprises an extension of the AFI/SAFI definitions to allow the address of the Next Hop for an IPv4 NLRI to also belong to the IPv6 protocol, the encoding of the Next Hop in order to determine which of the protocols the address actually belongs to, and a new BGP Capability allowing MP-BGP Peers to dynamically discover whether they can extend IPv4 NLRI with an IPv6 Next Hop.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

Table of Contents

1.	Introduction	3
2.	Extension of AFI/SAFI Definitions for IPv4 Address Family . .	4
3.	Use of BGP Capability Advertisement	5
4.	Operations	6
5.	Usage Examples	7
5.1.	IPv4 over IPv6 Core	7
5.2.	IPv4 VPN over IPv6 Core	7
6.	IANA Considerations	8
7.	Security Considerations	8
8.	Acknowledgments	8
9.	References	9
9.1.	Normative References	9
9.2.	Informative References	9
	Authors' Addresses	10
	Intellectual Property and Copyright Statements	11

1. Introduction

Multi Protocol-BGP (MP-BGP) ([[I-D.ietf-idr-rfc2858bis](#)]) specifies that the set of Network Layer protocols to which the address carried in the Next Hop field may belong is determined by the Address Family Identifier (AFI) and the Subsequent Address Family Identifier (SAFI). A number of existing AFI/SAFIs allow the Next Hop address to belong to a different address family than the Network Layer Reachability Information (NLRI). For example, the AFI/SAFI <25/65> used as per [[I-D.ietf-l2vpn-signaling](#)] in order to perform L2VPN Autodiscovery allows advertising an NLRI containing the identifier of a VPLS instance, or identifying a particular pool of attachment circuits at a given Provider Edge (PE), while the Next Hop field contains the loopback address of a PE. Similarly, the AFI/SAFI <1/132> defined in [[I-D.ietf-l3vpn-rt-constrain](#)] in order to advertise Route Target (RT) membership information, allows advertising an NLRI containing such RT membership information while the Next Hop field contains the address of the advertising router.

Furthermore, a number of these existing AFI/SAFIs allow the Next Hop to belong to either the IPv4 Network Layer Protocol or the IPv6 Network Layer Protocol, and specify the encoding of the Next Hop information in order to determine which of the protocols the address actually belongs to. For example, [[I-D.ietf-l3vpn-rt-constrain](#)] allows the Next Hop address to be either IPv4 or IPv6 and states that the Next Hop field address shall be interpreted as an IPv4 address, whenever the length of Next Hop address is 4 octets, and as a IPv6 address, whenever the length of the Next Hop address is 16 octets.

There are situations such as those described in [[I-D.ietf-softwire-problem-statement](#)] where carriers (or large enterprise networks acting as carrier for their internal resources) may be required to establish connectivity between 'islands' of networks of one address family type across a transit core of a differing address family type. This includes both the case of IPv6 islands across an IPv4 core and the case of IPv4 islands across an IPv6 core. Where Multi-Protocol BGP (MP-BGP) is used to advertise the corresponding reachability information, this translates into the requirement for a BGP speaker to advertise Network Layer Reachability Information (NLRI) of a given address family via a Next Hop of a different address family (i.e. IPv6 NLRI with IPv4 Next Hop and IPv4 NLRI with IPv6 Next Hop).

The current AFI/SAFI definitions for the IPv6 address family assume that the Next Hop address belongs to the IPv6 address family type. Specifically, as per [[RFC2545](#)] and [[RFC3107](#)], when the <AFI/SAFI> is <2/1>, <2/2>, or <2/4>, the Next Hop address is assumed to be of IPv6 type. As per [[RFC4659](#)], when the <AFI/SAFI> is <2/128>, the Next Hop

address is assumed to be of IPv6-VPN type.

However, [[I-D.ooms-v6ops-bgp-tunnel](#)] and [[RFC4659](#)] specify how an IPv4 address can be encoded inside the Next Hop IPv6 address field when an IPv6 NLRI needs to be advertised with an IPv4 Next Hop. [[I-D.ooms-v6ops-bgp-tunnel](#)] defines how the IPv4-mapped IPv6 address format specified in the IPv6 addressing architecture ([[RFC4291](#)]) can be used for that purpose when the <AFI/SAFI> is <2/1>, <2/2>, or <2/4>. [[RFC4659](#)] defines how the IPv4-mapped IPv6 address format as well as a null Route Distinguisher can be used for that purpose when the <AFI/SAFI> is <2/128>. Thus, there are existing solutions for the advertisement of IPv6 NLRI with an IPv4 next hop.

Similarly, the current AFI/SAFI definitions for the IPv4 address family assume that the Next Hop address belongs to the IPv4 address family type. Specifically, as per [MP-BGP] and [[RFC3107](#)], when the <AFI/SAFI> is <1/1>, <1/2>, or <1/4>, the Next Hop address is assumed to be of IPv4 type. As per [[RFC4364](#)], when the <AFI/SAFI> is <1/128>, the Next Hop address is assumed to be of VPN-IPv4 type. There is clearly no generally applicable method for encoding an IPv6 address inside the IPv4 address field of the Next Hop. Hence, there is currently no specified solution for advertising IPv4 NLRI with an IPv6 Next Hop.

This document specifies the extensions necessary to do so. This comprises an extension of the AFI/SAFI definitions to allow the address of the Next Hop for an IPv4 NLRI to belong to either the IPv4 or the IPv6 protocol, the encoding of the Next Hop information in order to determine which of the protocols the address actually belongs to, and a new BGP Capability allowing MP-BGP Peers to dynamically discover whether they can extend IPv4 NLRI with an IPv6 Next Hop. The new BGP capability allows gradual deployment of the new functionality of advertising IPv4 reachability via an IPv6 next hop, without any flag day nor any risk of traffic black-holing.

2. Extension of AFI/SAFI Definitions for IPv4 Address Family

As mentioned earlier, MP-BGP specifies that the set of Network Layer protocols to which the address carried in the Next Hop field may belong is determined by the Address Family Identifier (AFI) and the Subsequent Address Family Identifier (SAFI). The current AFI/SAFI definitions for the IPv4 address family only have provisions for advertising a Next Hop address that belongs to the IPv4 protocol. This document extends the definition of the AFI/SAFI associated with the AFI=1 (IPv4) to extend the set of network layer protocols to which the Next Hop address can belong, to include IPv6 in addition to IPv4.

Specifically, this document allows advertising with [\[I-D.ietf-idr-rfc2858bis\]](#) of an MP_REACH_NLRI with:

- o AFI=1
- o SAFI = 1, 2, 4 or 128
- o Length of Next Hop Address = 16 or 32
- o Next Hop Address = IPv6 address of next hop (potentially followed by the link-local IPv6 address of the next hop). This field is to be constructed as per [section 3 of \[RFC2545\]](#).
- o NLRI= NLRI as per current AFI/SAFI definition

This is in addition to the current mode of operation allowing advertisement of an NLRI for <AFI/SAFI> of <1,1>, <1,2> and <1,4> with a next hop address of IPv4 type and advertisement of an NLRI for <AFI/SAFI> of <1,128> with a next hop address of VPN-IPv4 type.

The BGP speaker receiving the advertisement MUST use the Length of Next Hop Address field to determine which network layer protocol the next hop address belongs to. When the Length of Next Hop Address field is equal to 16 or 32, the next hop address is of type IPv6. Note that this method of using the Length of the Next Hop Address field to determine which network layer protocol the next hop address belongs to (out of the set of protocols allowed by the AFI/SAFI definition) is the same as used in [\[I-D.ietf-l3vpn-rt-constrain\]](#).

[3.](#) Use of BGP Capability Advertisement

[\[RFC3392\]](#) defines a mechanism to allow two BGP speakers to discover if a particular capability is supported by their BGP peer and thus whether it can be used with that peer. This document defines a new capability which can be advertised using [\[RFC3392\]](#) and which is referred to as the Extended Next Hop Encoding capability. This capability allows BGP speakers to discover whether a peer supports advertisement of IPv4 NLRI with an IPv6 next Hop.

A BGP speaker that wishes to advertise to a BGP peer an IPv6 next hop for an IPv4 NLRI as per this specification MUST use the Capability Advertisement procedures defined in [\[RFC3392\]](#) to determine whether its peer supports this for the NLRI AFI/SAFI pair(s) of interest. The fields in the Capabilities Optional Parameter MUST be set as follows.

The Capability Code field MUST be set to [To be allocated by IANA]

(which indicates the Extended Next Hop Encoding capability).

The Capability Length field is set to a variable value which is the length of the Capability Value field (which follows).

The Capability Value field has the following format:

```
+-----+
| NLRI AFI - 1 (2 octets) |
+-----+
| NLRI SAFI - 1 (2 octets) |
+-----+
| Nexthop AFI - 1 (2 octets) |
+-----+
| ..... |
+-----+
| NLRI AFI - N (2 octets) |
+-----+
| NLRI SAFI - N (2 octet) |
+-----+
| Nexthop AFI - N (2 octets) |
+-----+
```

where each triple <NLRI AFI, NLRI SAFI, Nexthop AFI> indicates that an NLRI of NLRI AFI, NLRI SAFI may be advertised with a Next Hop address belonging to the Network Layer protocol of Nexthop AFI.

and where the AFI and SAFI values are defined in the Address Family Identifier and Subsequent Address Family Identifier registries maintained by IANA. Since this document concerns itself with the advertisement of IPv4 NLRI with an IPv6 Next Hop, this specification only defines the use of NLRI AFI=1 (IPv4) and Nexthop AFI=2 (IPv6). However, this encoding has been chosen to facilitate potential use of a similar mechanism in the future for other address families.

A BGP speaker MUST only advertise to a BGP peer an IPv4 NLRI with an IPv6 Next Hop if the BGP speaker has first ascertained via BGP Capability Advertisement that the BGP peer supports the Extended Next Hop Encoding capability for the relevant AFI/SAFI pair.

4. Operations

By default, if a particular BGP session is running over IPvx (where IPvx is IPv4 or IPv6), and if the BGP speaker sending an update is putting its own address in as the next hop, then the next hop address

SHOULD be specified as an IPvx address, using the encoding rules specified in the AFI/SAFI definition of the NLRI being updated. This default behavior may be overridden by policy.

When a next hop address needs to be passed along unchanged (as, e.g., a Route Reflector (RR) would do), its encoding MUST NOT be changed. If a particular RR client cannot handle that encoding (as determined by the BGP capability advertisement), then the NLRI in question cannot be distributed to that client. For sound routing in certain scenarios, this will require that all the RR clients be able to handle whatever encodings any of them may generate.

5. Usage Examples

5.1. IPv4 over IPv6 Core

The extensions defined in this document may be used for the interconnection of IPV4 islands over an IPv6 backbone. In this application, Address Family Border Routers (AFBR) (as defined in [[I-D.ietf-softwire-problem-statement](#)]) advertise IPv4 NLRI information in the MP_REACH_NLRI along with an IPv6 next hop.

The MP_REACH_NLRI is encoded with:

- o AFI=1
- o SAFI=1
- o Length of Next Hop Network Address = 16 (or 32)
- o Network Address of Next Hop= IPv6 address of Next Hop
- o NLRI= IPv4 routes

During BGP Capability Advertisement, the PE routers would include the following fields in the Capabilities Optional Parameter:

- o Capability Code set to "Extended Next Hop Encoding"
- o Capability Value containing <NLRI AFI=1, NLRI SAFI=1, Nexthop AFI=2>

5.2. IPv4 VPN over IPv6 Core

The extensions defined in this document may be used for support of IPV4 VPNs over an IPv6 backbone. In this application, PE Routers would advertise VPN-IPv4 NLRI information in the MP_REACH_NLRI along

with an IPv6 next hop.

The MP_REACH_NLRI is encoded with:

- o AFI=1
- o SAFI=128
- o Length of Next Hop Network Address = 16 (or 32)
- o Network Address of Next Hop= IPv6 address of Next Hop
- o NLRI= IPv4-VPN routes

During BGP Capability Advertisement, the PE routers would include the following fields in the Capabilities Optional Parameter :

- o Capability Code set to "Extended Next Hop Encoding"
- o Capability Value containing <NLRI AFI=1, NLRI SAFI=128, Nexthop AFI=2>

6. IANA Considerations

This document defines, in [section 3](#), a new Capability Code to indicate the Extended Next Hop Encoding capability in the [BGP-CAP] Capabilities Optional Parameter. The value for this new Capability Code is to be allocated from the range 1 through 63 set aside for allocation using the "IETF consensus" policy defined in [\[RFC2434\]](#).

7. Security Considerations

This document does not raise any additional security issues beyond those of BGP-4 and the MultiProtocol extensions for BGP-4. The same security mechanisms are applicable.

8. Acknowledgments

The authors would like to thank Yakov Rekhter for his contribution in the approach defined in this document.

9. References

9.1. Normative References

- [I-D.ietf-idr-rfc2858bis]
Bates, T., "Multiprotocol Extensions for BGP-4",
[draft-ietf-idr-rfc2858bis-10](#) (work in progress),
March 2006.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2545] Marques, P. and F. Dupont, "Use of BGP-4 Multiprotocol
Extensions for IPv6 Inter-Domain Routing", [RFC 2545](#),
March 1999.
- [RFC3107] Rekhter, Y. and E. Rosen, "Carrying Label Information in
BGP-4", [RFC 3107](#), May 2001.
- [RFC3392] Chandra, R. and J. Scudder, "Capabilities Advertisement
with BGP-4", [RFC 3392](#), November 2002.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing
Architecture", [RFC 4291](#), February 2006.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private
Networks (VPNs)", [RFC 4364](#), February 2006.

9.2. Informative References

- [I-D.ietf-l2vpn-signaling]
Rosen, E., "Provisioning, Autodiscovery, and Signaling in
L2VPNs", [draft-ietf-l2vpn-signaling-08](#) (work in progress),
May 2006.
- [I-D.ietf-l3vpn-rt-constrain]
Marques, P., "Constrained VPN Route Distribution",
[draft-ietf-l3vpn-rt-constrain-02](#) (work in progress),
June 2005.
- [I-D.ietf-softwire-problem-statement]
Li, X., "Softwire Problem Statement",
[draft-ietf-softwire-problem-statement-02](#) (work in
progress), May 2006.
- [I-D.ooms-v6ops-bgp-tunnel]
Clercq, J., "Connecting IPv6 Islands over IPv4 MPLS using
IPv6 Provider Edge Routers (6PE)",
[draft-ooms-v6ops-bgp-tunnel-06](#) (work in progress),
January 2006.

- [RFC2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 2434](#), October 1998.
- [RFC4659] De Clercq, J., Ooms, D., Carugi, M., and F. Le Faucheur, "BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN", [RFC 4659](#), September 2006.

Authors' Addresses

Francois Le Faucheur
Cisco Systems
Greenside, 400 Avenue de Roumanille
Sophia Antipolis 06410
France

Email: flefauch@cisco.com

Eric Rosen
Cisco Systems, Inc.
1414 Massachusetts Avenue
Boxborough, MA 01719
USA

Email: erosen@cisco.com

Full Copyright Statement

Copyright (C) The Internet Society (2006).

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 WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

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 [BCP 78](#) and [BCP 79](#).

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 implementers 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 that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

