

Network Working Group
Internet Draft
Expiration Date: October 2008

James Uttaro
AT&T

Pradosh Mohapatra
David J. Smith
Cisco Systems, Inc.

Robert Raszuk
John Scudder
Juniper Networks, Inc.

April 2008

BGP ACCEPT_OWN Well-known Community Attribute

[draft-pmohapat-idr-acceptown-community-01.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/lid-abstracts.txt>.

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

Abstract

It may be useful for a BGP speaker in an autonomous system to receive and accept its own advertised route from a route reflector with more fine-grained route control. For example, the route reflector can change certain attributes of a route as desired, and then re-

Internet Draftdraft-pmohapat-idr-acceptown-community-01.txt April 2008

advertise it back to the originator. Though it is possible to perform such policy control directly at the originator, it may be operationally cumbersome in a network with a large number of border routers having complex BGP policies.

This draft defines a new and well-known BGP community value, ACCEPT_OWN, that signals a BGP speaker to continue processing of an UPDATE message and the associated routes even when the ORIGINATOR_ID or the NEXT_HOP value matches that of the receiving speaker.

Table of Contents

| | | |
|--------------------|--|-------------------|
| 1 | Specification of Requirements | 2 |
| 2 | Introduction | 2 |
| 3 | ACCEPT_OWN Community | 3 |
| 4 | Security Considerations | 4 |
| 5 | IANA Considerations | 4 |
| 6 | Appendix A - Extranet application (non-normative) .. | 4 |
| 7 | Acknowledgements | 5 |
| 8 | Normative References | 5 |
| 9 | Informative References | 6 |
| 10 | Authors' Addresses | 6 |
| 11 | Full Copyright Statement | 7 |
| 12 | Intellectual Property | 7 |

[1](#). Specification of Requirements

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 [[RFC2119](#)].

[2](#). Introduction

In certain scenarios, a BGP speaker may maintain multiple contexts, in which case the speaker originates and receives routes within a

particular context (an example of such a context could be a VRF used by BGP/MPLS VPNs [[RFC4364](#)]). In such scenarios, the ability of a BGP speaker to accept a route with its own ORIGINATOR_ID or its own NEXT_HOP provides a way to modify and then redistribute routing information among the contexts maintained by the speaker through some

Internet Draftdraft-pmohapat-idr-acceptown-community-01.txt April 2008

other (external) speakers. For example, a route reflector can change certain path attributes of a route as desired, and then re-advertise it back to the originator. Though it is possible to perform such policy control directly on the originator, it may be operationally cumbersome in a network with a large number of border routers having complex BGP policies.

As per the BGP protocol [[RFC4271](#)], a BGP speaker rejects prefix advertisements received that were originated by itself. In an autonomous system with route reflectors, the route reflector attaches the ORIGINATOR_ID attribute to the UPDATE messages so that if such prefix advertisements reach the originator, the originator can reject them by simply checking the ORIGINATOR_ID attribute. The BGP specification also mandates that a route should not be advertised to a peer nor accepted from a peer when the NEXT_HOP attribute matches the receiver's own "IP address". These integrity checks help to detect and prevent routing information loops.

The draft proposes a modification to this behavior by defining a new well-known community [[RFC1997](#)] value. If this community value, ACCEPT_OWN, is attached to an UPDATE message, the originator will not reject the UPDATE message and the associated routes even when the ORIGINATOR_ID or the NEXT_HOP value matches that of the receiving speaker, thus enabling more fine-grained route control via a route reflector.

To prevent routing information loops, a BGP speaker SHOULD accept a route with its own ORIGINATOR_ID or NEXT_HOP value only if the ACCEPT_OWN community value is present and the context in which the speaker originated the route is different than the context in which the speaker accepts the route.

[3.](#) ACCEPT_OWN Community

This memo defines the use of a new well-known BGP non-transitive

community, ACCEPT_OWN, with value 0xFFFFF05. The ACCEPT_OWN community has global significance. However, it SHOULD NOT be advertised between external BGP peers. The ACCEPT_OWN community SHOULD only be advertised between internal BGP peers.

Use of this well-known community value signals that the associated route prefix should not be rejected by its originator irrespective of the ORIGINATOR_ID and NEXT_HOP values. The ACCEPT_OWN community effectively disables the ORIGINATOR_ID and NEXT_HOP integrity checks, however, only for those route prefixes having the ACCEPT_OWN community value.

Some route reflectors may be designed such that they never send routing information back to the router specified in ORIGINATOR_ID as mandated by [[RFC1966](#)]. Such route reflectors MUST disable this suppression functionality for routes which carry the ACCEPT_OWN community.

[4.](#) Security Considerations

ACCEPT_OWN as described above permits a router's own route prefix to be advertised to a different "context" on that router. In this respect, such a route is similar to any other BGP route and shares the same set of security vulnerabilities and concerns. No new fundamental security issues are introduced by ACCEPT_OWN.

[5.](#) IANA Considerations

This document defines a new well-known community, called ACCEPT_OWN. It is to be assigned value 0xFFFFF05.

[6.](#) [Appendix A](#) - Extranet application (non-normative)

One of the applications for this behavior is auto-configuration of extranets within MPLS VPN networks. Consider the following topology:

CE1 -----+

```

      |
      (VRF 1, RD 1, RT 1)
          PE1 ..... RR
      (VRF 2, RD 2, RT 2)
      |
CE2 -----+

```

Within the above topology, PE1 receives a prefix X from CE1. Prefix X is installed in VRF 1 and is advertised to the route reflector with route distinguisher (RD) 1 and route target (RT) 1 as configured on PE1. The requirement is to import prefix X into VRF 2 and advertise it to CE2 in support of extranet VPN connectivity between CE1/VRF1 and CE2/VRF2. Current BGP mechanisms for MPLS VPNs [[RFC4364](#)] require changing the import RT value and/or import policy for VRF 2 on PE1. This is operationally cumbersome in a network with a large number of border routers having complex BGP policies.

Alternatively, using the new ACCEPT_OWN community value, the route

reflector can simply re-advertise prefix X back to PE1 with RT 2 appended. In this way, PE1 will accept prefix X despite its ORIGINATOR_ID or NEXT_HOP value, import it into VRF 2, and will determine the correct adjacency rewrite within VRF 1 based on the RD value (1) and the prefix. The same operation needs also to happen in the reverse direction (VRF 1 learning a route from VRF 2) to achieve establishment of an extranet VPN strictly via the route reflector without changing the BGP policy of PE1 in any way.

A router performing such an extranet application can accept a route with its own ORIGINATOR_ID or NEXT_HOP value only if the "context" in which the router originated the route is different than the "context" in which the router accepts the re-advertised route (VRF is an example of a "context").

7. Acknowledgements

The authors would like to thank Yakov Rekhter, Jim Guichard, Clarence Filsfils, and John Mullooly for their valuable comments and suggestions.

8. Normative References

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

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

[RFC1997] Chandra, R., Traina, P., and T. Li, "BGP Communities Attribute", [RFC 1997](#), August 1996.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4364](#), February 2006.

[RFC1966] Bates, T. and Chandra, R, "BGP Route Reflection: An Alternative to full mesh IBGP," June 1996.

Uttaro, et al.

[Page 5]

Internet Draftdraft-pmohapat-idr-acceptown-community-01.txt April 2008

9. Informative References

[RFC3765] G. Huston, "NOPEER community for BGP route scope control", [RFC 3765](#), April 2004.

Schudel, G. and D. Smith, "Router Security Strategies: Securing IP Network Traffic Planes.", Cisco Press, January 2008.

10. Authors' Addresses

James Uttaro
AT&T
200 S. Laurel Avenue

Middletown, NJ 07748
Email: uttaro@att.com

Pradosh Mohapatra
Cisco Systems, Inc.
170 Tasman Drive
San Jose, CA 95134
Email: pmohapat@cisco.com

David J. Smith
Cisco Systems, Inc.
499 Thornall Street
Edison, NJ 08837
E-mail: dasmith@cisco.com

Robert Raszuk
Juniper Networks
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
Email: raszuk@juniper.net

John Scudder
Juniper Networks
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
Email: jgs@juniper.net

11. Full Copyright Statement

Copyright (C) The IETF Trust (2008).

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, THE IETF TRUST 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.

12. 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

ipr@ietf.org.