

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 9, 2017

K. Patel
A. Lindem
Cisco Systems
L. Jalil
Verizon
July 8, 2016

Selective Advertisement of Multiple Paths within BGP
draft-keyupate-idr-bgp-selective-add-paths-01.txt

Abstract

[[draft-ietf-idr-add-paths](#)] defines a BGP extension that allows the advertisement of multiple paths for the same address prefix without the new paths implicitly replacing any previous ones. The essence of the extension is that each path is identified by a path identifier in addition to the address prefix. This draft augments functionality defined in [[draft-ietf-idr-add-paths](#)] to facilitate advertisement of multiple paths for a subset of prefixes in a given address family. Prefixes are selected through specification of a well-known BGP extended community.

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 January 9, 2017.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of

publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	2
1.1.	Requirements Language	3
2.	Selective Add-Path Capability	3
3.	Selective Add-Path Community	4
4.	Selective Add-Path Use Case	5
5.	IANA Considerations	5
6.	Security Considerations	5
6.1.	Acknowledgements	5
7.	References	5
7.1.	Normative References	6
7.2.	Information References	6
	Authors' Addresses	6


[1.](#) Introduction

[I-D.ietf-idr-add-paths] defines a BGP extension that allows the advertisement of multiple paths for the same address prefix without the new paths implicitly replacing any previous ones. The essence of the extension is that each path is identified by a path identifier in addition to the address prefix. This document augments functionality defined in defined in [\[I-D.ietf-idr-add-paths\]](#) to facilitate advertisement of multiple paths for a subset of prefixes in a given

address family. Prefixes are selected through specification of a reserved BGP extended community.

This draft defines a capability to limit the scope of BGP multiple path advertisement to a subset prefixes in a given address family.

Prefixes are selected through specification of a reserved BGP extended community [[RFC4360](#)].



As an example, suppose that RR is a route reflector that doesn't change nexthops of the prefixes it reflects, with clients R1, R2 and R3. Suppose R1 sends RR an UPDATE: <NLRI=P1, NH=R1> and <NLRI=P2, NH=R1>. Suppose R2 sends RR an UPDATE: <NLRI=P1, NH=R2> and <NLRI=P2, NH=R2>. R1, R2, and R3 would like selective ADDPATHs for Prefix P1 and not for Prefix P2. R1, R2, and R3 exchange selective the ADDPATH capability with RR. R1, R2, R3 are configured with the reserved selective ADDPATHs community that they attach to prefixes that need selective ADDPATHs. RR now has two paths to P1 and P2. RR announces P2 with bestpath to all its clients while RR announces P1 with additional paths. The number of additional paths with its best path and its additional paths is a matter of local policy configured on RR.

[1.1.](#) 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](#)].

[2.](#) Selective Add-Path Capability

The ADD-PATH Capability is a new BGP capability [[RFC5492](#)]. The Capability Code for this capability is allocated by IANA as specified in the [Section 5](#). The Capability Length field of this capability is variable. The Capability Value field consists of one or more of the following tuples:

```
+-----+
| Address Family Identifier (2 octets)      |
+-----+
| Subsequent Address Family Identifier (1 octet) |
+-----+
```

The meaning and use of the fields are as follows:

Address Family Identifier (AFI):

This field is the same as the one used in [[RFC4760](#)].

Subsequent Address Family Identifier (SAFI):

This field is the same as the one used in [[RFC4760](#)].

A BGP Speaker that wishes to announce or receive multiple paths MUST exchange the add-path capability defined in [[I-D.ietf-idr-add-paths](#)]. A BGP Speaker that wishes to announce or receive multiple paths for selected prefixes MUST exchange the selective add-path capability defined in this draft. A BGP speaker wanting to advertise selective add-path capability MUST also advertise the add-path capability defined in [[I-D.ietf-idr-add-paths](#)].

In processing a received selective add-path capability from a peer, a BGP speaker MUST ensure that it also received the add-path capability defined in [[I-D.ietf-idr-add-paths](#)]. Otherwise, the BGP speaker

should ignore the received selective add-path capability and follow the error handling rules for unsupported add-path capabilities in [\[RFC5492\]](#).

[3.](#) Selective Add-Path Community

Upon successful Selective Add-Path capability negotiation, a BGP speaker MUST NOT announce multiple paths for any AFI/SAFI prefix unless it has received at least one UPDATE for that prefix that includes the Selective Add-Path well-known community in its attributes. The community is a Transitive Opaque Extended Community with the sub-type value IANA-TBD.

If Selective Add-Path capability negotiation for a given AFI/SAFI has not taken place and the Selective Add-Path Community is included with a prefix advertised for the same AFI/SAFI, the Selective Add-Path Community will be ignored. However, the occurrence of the unexpected community SHOULD be logged.

[4.](#) Selective Add-Path Use Case

A use case is a BGP deployment where underlay and overlay routes are associated with the same AFI/SAFI and, due to scaling, only multiple paths are only advertised and installed for underlay routes. For direct BGP sessions, the ingress routers would only advertise multiple paths for the underlay routes. However, if the topology includes BGP Router Reflectors [\[RFC4456\]](#), it is likely that multiple ingress routers will advertise the same overlay routes. In this case, the mechanism describe herein would be useful in limiting multi-path best-path computation and advertisement to the underlay routes.

As a second usecase, many times a service provider will carry both customer traffic and internal services (e.g., VOIP) on the same backbone network using routes in the same BGP address families. In this situation, the number of customer routes and paths greatly exceed the number of routes and paths for internal services. However, the service provider desires the faster failover and convergence provided by BGP Add-Paths [\[I-D.ietf-idr-add-paths\]](#). In this scenario, the Selective Add-Path functionality described herein can be leveraged for routes corresponding to internal services

without the overhead incurred if multiple paths were advertised for the customer routes.

[5.](#) IANA Considerations

This document defines a new capability for BGP. We request IANA to assign BGP capability number from BGP Capabilities Registry.

This document also defines a new extended community for BGP. We request IANA to assign a BGP well-known extended community from the Transitive Opaque Extended Community Sub-Types Registry.

[6.](#) Security Considerations

This extension to BGP does not change the underlying security issues inherent in the existing [[RFC4724](#)] and [[RFC4271](#)].

[6.1.](#) Acknowledgements

The authors would like to thank for the review and comments.

[7.](#) References

Patel, et al. Expires January 9, 2017 [Page 5]

Internet-Draft BGP Add-Path Selective Advertisement July 2016

[7.1.](#) Normative References

- [I-D.ietf-idr-add-paths]
Walton, D., Retana, A., Chen, E., and J. Scudder,
"Advertisement of Multiple Paths in BGP", [draft-ietf-idr-add-paths-13](#) (work in progress), December 2015.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), DOI 10.17487/RFC4271, January 2006,

<<http://www.rfc-editor.org/info/rfc4271>>.

- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", [RFC 4360](#), DOI 10.17487/RFC4360, February 2006, <<http://www.rfc-editor.org/info/rfc4360>>.
- [RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", [RFC 5492](#), DOI 10.17487/RFC5492, February 2009, <<http://www.rfc-editor.org/info/rfc5492>>.

7.2. Information References

- [RFC4456] Bates, T., Chen, E., and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)", [RFC 4456](#), DOI 10.17487/RFC4456, April 2006, <<http://www.rfc-editor.org/info/rfc4456>>.
- [RFC4724] Sangli, S., Chen, E., Fernando, R., Scudder, J., and Y. Rekhter, "Graceful Restart Mechanism for BGP", [RFC 4724](#), DOI 10.17487/RFC4724, January 2007, <<http://www.rfc-editor.org/info/rfc4724>>.

Authors' Addresses

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

Email: keyupate@cisco.com

Patel, et al.

Expires January 9, 2017

[Page 6]

Internet-Draft

BGP Add-Path Selective Advertisement

July 2016

Acee Lindem
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
USA

Email: acee@cisco.com

Luay Jalil
Verizon
400 International Parkway
Richardson, Tx 75081
USA

Email: luay.jalil@verizon.com