Global Routing Operations
Internet-Draft
Updates: 7947,7948 (if approved)
Intended status: Informational
Expires: September 12, 2019

M. Aelmans
Juniper Networks
S. Konstantaras
AMS-IX
S. Plug
ECIX - Megaport
C. Dietzel
DE-CIX

March 11, 2019

BGP Large Communities applications for IXP Route Servers
draft-adkp-grow-ixpcommunities-00

Abstract

This document presents suggestion and examples for application of BGP Large Communities [RFC8092] at Internet Exchange Points (IXPs). Suggestions are based on operational experiences from IXP operators and members. Any IXP operator or IXP member can consider using these communities. The document specifically focusses on Route Server [RFC7947] deployments in IXP context [RFC7948].

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

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 https://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 12, 2019.

Table of Contents

1. Introduction

This document presents suggestions for the application of BGP Large Communities [RFC8092] to IXP operators and members using the BGP [RFC4271] protocol. It adds specific suggestions for the operators and members of IXPs deploying BGP Large Communities as suggested in [RFC8195].

2. Justification

Networks operating in the DFZ tend to exchange routing information at multiple IXP in order to improve redundancy and geographical optimization. Besides 'the typical' IXP members an increasing amount of enterprise networks connect to IXPs. They have additional requirements. In order to offer a uniform mode of operation across different IXPs there is a need to provide standards.

3. Suggested Large BGP Community Standard List

This list proposes a standard to use in IXP operations for the use of BGP Large Communities. It was first published at the EURO-IX website [EURO-IX].

The tables below provide a per 'section' divided overview of Large Community usage.

Range	Description	Notes	Strip on export	Priority
RS:0:PEERAS	Do not advertise to PEERAS		recommended	0
RS:1:PEERAS	Advertise to PEERAS	Only useful in combination with RS:0:0	recommended	1
RS:2:ms	Do not announce to peers higher than ms	ms = Latency of peer in ms	recommended	2

Table 1: Direct filtering RS:0-99:*

Range	Description	Notes	Strip on export	Priority
RS:101:PEERAS	Prepend to PEERAS once		yes	3
RS:102:PEERAS	Prepend to PEERAS twice		yes	3
RS:103:PEERAS	Prepend to PEERAS three times		yes	3
RS:111:ms	Prepend once to peers higher than ms		yes	3
RS:112:ms	Prepend twice to peers higher than ms		yes	3
RS:113:ms	Prepend three to peers higher than ms		yes	3

Table 2: AS Path prepending RS:100-199:*

Range	Description	Notes	Strip on export	Priority

Table 3: Unassigned RS:200-899:*

Range	Description	Notes	Strip on export	Priority

Table 4: Informational RS:1000-1999:*

Range	Description	Notes	Strip on export

RS:1000:1	RPKI VALID	Prefix is RPKI VALID	yes
RS:1000:2	RPKI UNKNOWN	Prefix is RPKI UNKNOWN	yes
RS:1000:3	RPKI NOT CHECKED		yes
RS:1000:4-*	Prefix is RPKI INVALID because of $REASON		yes
RS:1001:1	IRRDB VALID	Prefix exists in IRRDB	yes
RS:1001:2	IRRDB NOT CHECKED	Prefix was not checked in IRRDB	yes
RS:1001:3	MORE SPECIFIC THAN IRRDB	Prefix does not exist in IRRDB, but a less specific does valid entry exists	yes
RS:1001:4	IRRDB Prefix not found in AS-SET or aut-num	Prefix was not found in the peer's as-set	yes
RS:1001:5	IRRDB INVALID ORIGIN AS	Origin AS not in peer AS-SET	yes
RS:1001:6	IRRDB INVALID PREFIX FOR ORIGIN AS	Prefix not found in origin AS	yes
RS:1002:1-*	TRACER (RS #)	IXP assigned ID for route server instance	no
RS:1003:ms	Measured RTT for advertising peer	IXP measured round trip time for peer in ms	yes
RS:1004:$peerAS	Incoming Peer AS	Use Autonomous System Number of the incoming	yes

		member for that route	
RS:1005:1	AS Object, Route Object and Organization NOT from the same region	Meant as a transitioning mechanism until full RPKI deployment	yes
RS:1005:2	AS Object, Route Object and Organization from within the same region	yes	
RS:1005:3	AS Object, Route Object and Organization from within the same region Not checked	yes	

Table 5: Informational tags RS:1000-1099:*

Range	Description	Notes	Strip on export
RS:1101:1	Prefix length too long		
RS:1101:2	Prefix length too short		
RS:1101:3	Bogon Prefix		
RS:1101:4	Bogon AS		
RS:1101:5	AS path too long		
RS:1101:6	AS path too short		
RS:1101:7	as-path.first != peeras		
RS:1101:8	next hop IP != peer IP		

RS:1101:9	IRRDB Prefix not found in AS-SET or aut-num	Prefix was not found in the peer's as-set	
RS:1101:10	Origin AS not in peer AS-SET		
RS:1101:11	Prefix not found in origin AS		
RS:1101:12	Prefix is RPKI UNKNOWN		
RS:1101:13	Prefix is RPKI INVALID		
RS:1101:14	Transit-free ASN in AS-Path		
RS:1101:15	Too many BGP communities set on prefix		

Table 6: Informational RS:1000-1999:*

Range	Description	Notes	Strip on export
RS:1102:1	Advertising peer declines prefix	Advertising peer does not want you to receive prefix	
RS:1102:2	You declined prefix from advertising peer	You do not want to receive prefix from advertising peer	
RS:1102:3	Maximum number of BGP communities exceeded		

Table 7: Route was filtered on export RS:1102:*

Range	Description	Notes	Strip on export

Table 8: Unassigned RS:1200-1899:*

Range	Description	Notes	Strip on export

Free to use informational communities

Table 9: IXP Specific RS:1900-1999:*

4. Security Considerations

Operators should note the recommendations in Section 11 of BGP Operations and Security [RFC7454] and handle BGP Large Communities with their ASN in the Global Administrator field similarly.

5. IANA Considerations

6. Acknowledgments

The authors would like to thank Colby Barth (Juniper Networks) and Bijal Sanghani (EURO-IX) for their support, insightful review, and comments.

7. Appendix: Implementation Guidance

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://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, <https://www.rfc-editor.org/info/rfc4271>.

[RFC7454] Durand, J., Pepelnjak, I., and G. Doering, "BGP Operations and Security", BCP 194, RFC 7454, DOI 10.17487/RFC7454, February 2015, <https://www.rfc-editor.org/info/rfc7454>.

[RFC7947] Jasinska, E., Hilliard, N., Raszuk, R., and N. Bakker, "Internet Exchange BGP Route Server", RFC 7947, DOI 10.17487/RFC7947, September 2016, <https://www.rfc-editor.org/info/rfc7947>.

[RFC7948] Hilliard, N., Jasinska, E., Raszuk, R., and N. Bakker, "Internet Exchange BGP Route Server Operations", RFC 7948, DOI 10.17487/RFC7948, September 2016, <https://www.rfc-editor.org/info/rfc7948>.

[RFC8092] Heitz, J., Ed., Snijders, J., Ed., Patel, K., Bagdonas, I., and N. Hilliard, "BGP Large Communities Attribute", RFC 8092, DOI 10.17487/RFC8092, February 2017, <https://www.rfc-editor.org/info/rfc8092>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

8.2. Informative References

[EURO-IX] EURO-IX, "Large BGP Communities", April 2018, <https://www.euro-ix.net/en/forixps/large-bgp-communities/>.

[RFC8195] Snijders, J., Heasley, J., and M. Schmidt, "Use of BGP Large Communities", RFC 8195, DOI 10.17487/RFC8195, June 2017, <https://www.rfc-editor.org/info/rfc8195>.

Authors' Addresses

Melchior Aelmans
Juniper Networks
Boeing Avenue 240
Schiphol-Rijk 1119 PZ
The Netherlands

Email: maelmans@juniper.net

Stavros Konstantaras
AMS-IX
Frederiksplein 42
Amsterdam 1017 XN
The Netherlands

Email: stavros.konstantaras@ams-ix.net

Stefan Plug
ECIX - Megaport
Tauentzienstr. 11
Berlin 10789
Germany

Email: spl@ecix.net

Christoph Dietzel
DE-CIX
Lindleystr. 12
Frankfurt am Main 60314
Germany

Email: christoph.dietzel@de-cix.net