Global Routing Operations
Internet-Draft
Intended status: Informational
Expires: August 20, 2017

J. Snijders
J. Heasley
NTT
M. Schmidt
i3D.net
February 16, 2017

Usage of BGP Large Communities
draft-ietf-grow-large-communities-usage-02

Abstract

Examples and inspiration for operators for the use of BGP Large Communities.

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 [RFC2119].

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 August 20, 2017.

Copyright Notice

Table of Contents

1. Introduction

BGP Large Communities [RFC8092] provide a mechanism to signal opaque information between Autonomous Systems. This document presents a set of examples of how Large BGP Communities could be employed by an operator to achieve various goals. This document draws from experience in Operational Communities such as NANOG [1] and NLNOG [2].

The opaque nature of BGP Large Communities allows for rapid deployment of new features or changes to the product. Operators are encouraged to publicly publish and maintain documentation of the purpose of each Large BGP Community, both informational and action, that they support or are visible in looking glasses.

2. The Generic Design Pattern

BGP Large Communities are composed of a 4-octet Global Administrator field followed by two 4-octet Local Data fields. Large BGP Communities are compose three 4-octet fields. The first is the Global Administrator field, whose value is the ASN of AS that has defined the meaning of the remaining two 4-octet fields, the Local Data fields. This document describes an approach defining these fields as "ASN:Function:Parameter"-approach to fill the three fields.

In deployments of both BGP Communities [RFC1997] and BGP Large Communities, two categories of Communities exist:

- Informational Communities
- Action Communities

For each, ideas are provided regarding the contents of each of the three fields in BGP Large Communities.

Throughout the document a topology of four Autonomous Systems is used to illustrate the usage of Communities in the following configuration:

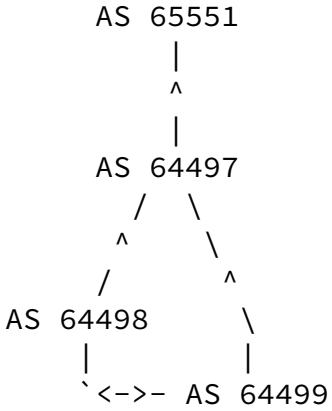

AS 64497 obtains transit services from AS 65551, a 32-bit ASN. AS 64497 provides transit services to both AS 64498 and AS 64499. AS 64498 and AS 64499 maintain a peering relationship in which they only exchange their customer routes.

2.1. Informational Communities

Informational Communites are labels for attributes such as origin of the route announcement, the relation with the EBGP neighbor or for instance the intended propagation audience. Informational Communities also assist in network operations such as debugging.

The Global Administrator field is set to the ASN which is marking the routes with the Informational Communities. For example, AS 64497 might add a community with the GA 64497 to a route learned from an iBGP or eBGP peer that means that the route was learned from or originated by a device in the Netherlands.

In general the intended audience of Informational Communities are downstream networks and the Global Administrator itself, but any Autonomous System could benefit from receiving these communities.

2.2. Action Communities

Action Communities are attached to routes to request non-default behaviour in this, a conferation or an external AS. Action Communities could be used to change the route's propagation characteristics, the LOCAL_PREFENCE or the number of AS_PATH prepends to add when exporting or importing a route.

The Global Administrator field is set to the ASN which is expected to perform the action. For instance, AS 64499 might add a Large Community with the GA 64497 to signal AS 64497 to perform an action upon that route.

In general the intended audience of Action Communities is an upstream provider, but realistically could be any AS willing to act upon it.

3. Examples of Informational Communities

3.1. Location

AS 64497 can inform its downstream networks about the geographical entity where AS 64497 learned a route by marking the route with BGP Large Communities following one or a combination of the following schemes.

3.1.1. An ISO 3166-1 numeric function

AS 64497 could assign a value of 1 to the first Local Data field to designate the function of the second Local Data field as ISO-3166-1 numeric country identifiers.

BGP Large Community	Description
64497:1:528	Route learned in Netherlands
64497:1:392	Route learned in Japan
64497:1:840	Route learned in United States of America

Example documentation for AS 64497 using Informational Communities describing the origin of routes using ISO 3166-1 numeric identifiers.

Table 1: Information: ISO 3166-1

3.1.2. An UNSD region function

AS 64497 could assign a value of 2 to the first Local Data field to designate the function of the parameter in the second Local Data field as an identifier for the macro geographical (continental) regions, geographical sub-regions, or selected economic and other groupings following a set of published identifiers by the United Nations Statistics Division [3].

BGP Large Community	Description
64497:2:2	Route learned in Africa
64497:2:9	Route learned in Oceania
64497:2:145	Route learned in Western Asia
64497:2:150	Route learned in Europe

Example documentation for AS 64497 using Informational Communities describing the origin of routes using numeric identifiers provided by the UN Statistics Division.

Table 2: Information: Regions

3.2. Relation

AS 64497 could assign a value of 3 to the first Local Data field to designate that the second Local Data field contains an identifier showing the relation with the EBGP neighbor from whom the route was received.

BGP Large Community	Description
64497:3:1	Route learned from a customer
64497:3:2	Route learned from a peering partner
64497:3:3	Route learned from an upstream provider

Example documentation for AS 64497 using Informational Communities describing the relation with the ASN from which the route was received.

Table 3: Information: Relation

3.3. Combining Informational Communities

Multiple Informational Communities can be tagged on a route, for example: a route learned in the Netherlands from a customer can contain both 64497:1:528 and 64497:2:150 and 64497:3:1.

4. Examples of Action Communities

4.1. Selective NO_EXPORT

As part of the commercial agreement between AS 64497 and AS 64498, AS 64497 might offer AS 64498 certain BGP Traffic Engineering features such as selectively not export routes learned from 64498 to certain EBGP neighbors of AS 64497.

4.1.1. Peer ASN Based Selective NO_EXPORT

AS 64497 might assign function identifier 4 to allow preventing propagation of routes to the ASN listed in the second Local Data field.

BGP Large Community	Description
64497:4:2914	Do not export route to AS 2914
64497:4:7018	Do not export route to AS 7018
64497:4:65551	Do not export route to AS 65551

Example documentation for AS 64497 offering Action Communities to limit propagation of routes based on the Peer ASN described in the third field.

Table 4: Action: Peer ASN NO_EXPORT

4.1.2. Location Based Selective NO_EXPORT

AS 64497 might assign function identifier 5 to allow its customers to request selectively not exporting routes on EBGP sessions within a certain geographical area. This example follows the ISO 3166-1 numeric encoding.

BGP Large Community	Description
64497:5:528	Do not export to EBGP neighbors in the Netherlands
64497:5:392	Do not export to EBGP neighbors in Japan
64497:5:840	Do not export to EBGP neighbors in United States of America

Example documentation for AS 64497 offering Action Communities to trigger NO_EXPORT on routes only when propagating the route to a certain geographical region.

Table 5: Action: NO_EXPORT in Region

4.2. Selective AS_PATH Prepending

As part of the commercial agreement between AS 64497 and AS 64498, AS 64497 might offer AS 64498 certain BGP Traffic Engineering features such as selectively prepending the AS_PATH with 64497's ASN to certain EBGP neighbors of AS 64497.

4.2.1. Peer ASN Based Selective AS_PATH Prepending

AS 64497 might assign function identifier 6 to allow prepending the AS_PATH on propagation of routes to the ASN listed in the second Local Data field.

BGP Large Community	Description
64497:6:2914	Prepend 64497 once on export to AS 2914
64497:6:7018	Prepend 64497 once on export to AS 7018
64497:6:65551	Prepend 64497 once on export to AS 65551

Example documentation for AS 64497 offering Action Communities to trigger prepending of the AS_PATH only when propagating the route to a certain Peer ASN.

Table 6: Action: Prepend to Peer ASN

4.2.2. Location Based Selective AS_PATH Prepending

AS 64497 might assign function identifier 7 to allow prepending of the AS_PATH on propagation of routes to on any EBGP neighbor's interconnection in the geographical entity listed in the second Local Data field. This example follows the ISO 3166-1 numeric regions codes in the Local Data 2 field.

BGP Large Community	Description
64497:7:528	Prepend once to EBGP neighbors in the Netherlands
64497:7:392	Prepend once to EBGP neighbors in Japan
64497:7:840	Prepend once to EBGP neighbors in United States of America

Example documentation for AS 64497 offering Action Communities to trigger prepending of the AS_PATH only when propagating the route to a certain geographical region.

Table 7: Action: Prepend in Region

4.3. Location based manipulation of LOCAL_PREF

In some cases, it can be desirable for an autonomous system to allow adjacent Autonomous Systems to directly influence the degree of preference associated with a route, usually expressed within the LOCAL_PREF attribute.

Furthermore, in the case of large networks spanning significant geography, it is often also useful to be able to extend this

capability and scope its effect to a geographic region. This is a more powerful mechanism than AS_PATH prepending, but since degree of preference determines BGP route selection and thus onward advertisement, it can also be self-limiting in its scope.

Since the LOCAL_PREF attribute which influences degree of preference is locally significant within each autonomous system, it is not usually practical or useful to compare LOCAL_PREF attribute values between autonomous systems. Instead it can be useful to classify the major types of route likely to exist within an autonomous system's routing hierarchy and provide an ability to set one's route to that preference:

- A qualified customer route. Usually the highest preference.
- A peer, or network-share, route. A co-operating network provider engaged in a partnership for customer coverage ("peering").
- A last resort, or backup route.

It is entirely possible that some providers may have more classes of route preference but it is possible to codify both the route preference class and the regional scope within the Local Data fields of the Large Community attribute.

For example, AS64497 might establish the following function identifiers to set route preference class, which could allow pairing with a location or peer-based operand to determine scope.

Function	Preference Class
10	Qualified customer route. Highest preference.
15	Peering partner. Median preference.
19	Route of last resort. Lowest preference.

Table 8: Action: Preference Function Identifiers

Once established, these route preference setting functions can be linked with a scoping operand such as per-peer or per-location based identifiers in order to provide AS64497's customers with a comprehensive and rich toolset to influence route preference.

BGP Large Community	Description
64497:15:528	Set as peer route in Netherlands
64497:19:840	Set as backup route in United States of America

Table 9: Action: Regional Preference Communities

Since the degree of preference influences BGP best path selection (which in turn influences onward route propagation) Operators should take special care with a traffic engineering tool such as location based local preference influence (BGP Wedgies [RFC4264]).

5. Security Considerations

Network operators should note the recommendations in Section 11 of BGP Operations and Security [RFC7454].

6. IANA Considerations

None.

7. Acknowledgements

The authors would like to gratefully acknowledge the insightful comments, contributions, critique and support from John Heasley, Adam Chappell, Jonathan Stewart, and Will Hargrave.

8. References

8.1. Normative References

[RFC1997] Chandra, R., Traina, P., and T. Li, "BGP Communities Attribute", RFC 1997, DOI 10.17487/RFC1997, August 1996, <http://www.rfc-editor.org/info/rfc1997>.

[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>.

[RFC4264] Griffin, T. and G. Huston, "BGP Wedgies", RFC 4264, DOI 10.17487/RFC4264, November 2005, <http://www.rfc-editor.org/info/rfc4264>.

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

[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, <http://www.rfc-editor.org/info/rfc8092>.

8.2. URIs

[1] http://nanog.net

[2] http://nlnog.net

[3] http://unstats.un.org/unsd/methods/m49/m49regin.htm

Authors' Addresses

Job Snijders
NTT Communications
Theodorus Majofskistraat 100
Amsterdam 1065 SZ
The Netherlands

Email: job@ntt.net

John Heasley
NTT Communications
12160 NW Coleman Drive
Portland, OR 97229

United States of America

Email: heas@shrubbery.net

Martijn Schmidt
i3D.net
Rivium 1e Straat 1
Capelle aan den IJssel 2909 LE
NL

Email: martijnschmidt@i3d.net