

Workgroup: IDR
Internet-Draft: draft-heitz-idr-wklc-05
Published: 6 September 2022
Intended Status: Standards Track
Expires: 10 March 2023
Authors: J. Heitz K. Sriram B. Dickson J. Heasley
 Cisco NIST
 BGP Well Known Large Community

Abstract

A range of BGP Autonomous System Numbers is reserved to create a set of BGP Well Known 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 <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 10 March 2023.

Copyright Notice

Copyright (c) 2022 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 (<https://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 Revised BSD License text as described in

Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

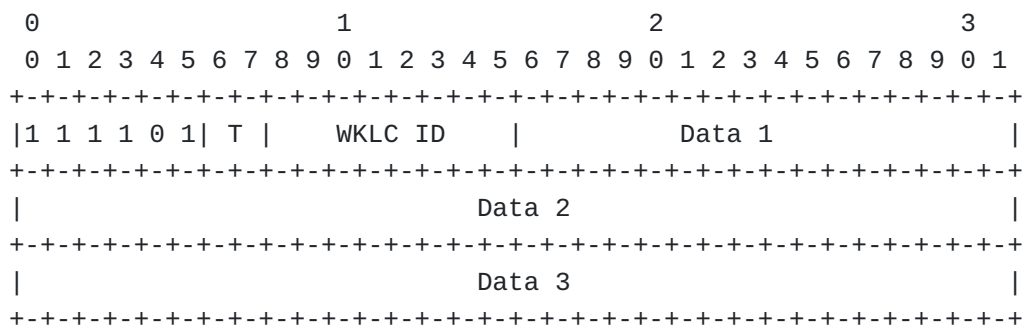
1. [Introduction](#)
 2. [Encoding](#)
 3. [Transitivity](#)
 4. [Discussion](#)
 5. [Security Considerations](#)
 6. [IANA Considerations](#)
 7. [Normative References](#)
- [Authors' Addresses](#)

1. Introduction

The Global Administrator field of the BGP Large community [[RFC8092](#)] is an Autonomous System Number (ASN). To create a set of Well Known Large Communities, a set of ASNs must be reserved for them, such that a real ASN in the Global Administrator field cannot be mistaken for a Well Known Large Community.

2. Encoding

Each BGP Well Known Large Community value is encoded as a 12-octet quantity, as follows:



The fields are as shown below:

T - Transitivity field (2 bits). This is further described below.

WKLC ID - Well Known Large Community Identifier (1 octet). See [IANA Considerations](#). If an experimental type is used, then it MUST NOT be hard coded in the BGP speaker software; it MUST be configurable. Different experiments can then run in the same network without having to coordinate identifier assignment during the coding stage.

Data 1,2,3 - A 10 octet value specific to the WKLC. Data 1 is 16 bits long and Data 2 and Data 3 are 32 bits long each. The data

portion is divided into the three fields only to encourage a canonical representation that follows [[RFC8092](#)]. If any WKLC cannot make use of this data division, it is free to define another division.

3. Transitivity

The Large Community path attribute is a transitive attribute. Thus, BGP speakers that do not implement the transitivity described here will transit the WKLC regardless. If such a speaker wishes not to receive a particular large community, it MUST filter it out using local policy. The transitivity field determines how BGP speakers transfer the WKLC across real Autonomous System (AS) boundaries. The values are:

- 0 - Transitive: The WKLC is transitive across ASes.
- 1 - Non-transitive: The WKLC is not transitive across ASes.
- 2 - Administration Transitive: The WKLC is transitive across ASes under the same administration only. By default, every AS boundary is also an administration boundary. If an external BGP session is configured as a non-administrative boundary, then it will send and receive WKLCs with transitivity 2, else it will discard the WKLC from the UPDATE message.
- 3 - One-time Transitive: The WKLC is transitive across ASes under the same administration and into an AS under the neighboring administration, but not into an AS under a further administration. A BGP speaker that receives a WKLC with transitivity 3 on an external BGP session on an administrative boundary SHOULD change the transitivity to 2.

4. Discussion

A criterion considered important is the number of data octets available for any WKLC type. This is maximized to 10 at the expense of ASN space and type space.

8 bits for the type is considered to be plenty. 255 types are more than enough for WKLC, considering how many extended community types have been used so far (22 at time of this publication). If a large set of types, each of which requires less than 10 octets of data, is required, then they can all be specified under a single type code and further distinguished by using the Data 1 field as a sub-type. An equivalent example is the EVPN Extended Community type, which defines further sub types.

The range of AS numbers currently unallocated by IANA is 399,261 to 4,199,999,999. The WKLC reserves 67,108,864 AS numbers. That still

leaves 4,132,491,874 unallocated AS numbers. For comparison, there are 94,968,317 AS numbers reserved for private use. Thus, the number of ASNs reserved for WKLCs is considered insignificant.

5. Security Considerations

The BGP Large Community Path attribute is transitive. Thus, a BGP speaker that does not recognize the transitivity field may transmit the WKLC contrary to the advisement of the transitivity field. If a BGP speaker wishes not to receive any Large Community, it must continue to filter it in the same way it was doing before the transitivity field was introduced.

In some cases, a received route that contains an AS number from the range assigned to WKLC in its AS_PATH may be considered unusual. A receiving BGP speaker MAY drop such a received route in route-policy.

6. IANA Considerations

IANA is requested to assign the range 4093640704 (0xF4000000) to 4160749567 (0xF7FFFFFF) from the BGP ASN registry for BGP Well Known Large Communities.

IANA is requested to create a registry of Well Known Large Communities in the range 0 to 255. Numbers from this registry are to be assigned in accordance with the policies defined in [RFC8126]. The policies for the following number ranges are:

0-63 - RFC Required

64-223 - First Come First Served

224-255 - Experimental

7. 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>>.
- [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>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26,

Authors' Addresses

Jakob Heitz
Cisco
170 West Tasman Drive
San Jose, CA 95134
United States of America

Email: jheitz@cisco.com

Kotikalapudi Sriram
USA National Institute of Standards and Technology
100 Bureau Drive
Gaithersburg, MD 20899
United States of America

Email: kotikalapudi.sriram@nist.gov

Brian Dickson

Email: brian.peter.dickson@gmail.com

John Heasley

Email: heas@shrubbery.net