

BGP Large Communities Attribute

RFC 8092

Greg Hankins

greg.hankins@nokia.com

Nokia

A Brief History of BGP Communities

- BGP Communities Attribute ([RFC 1997](#), August 1996)
 - Designed to simplify Internet routing policies
 - Encodes a 32-bit value displayed as “16-bit ASN:16-bit value”
 - Broad support in BGP implementations, and widely deployed by network operators for Internet routing
 - For example: 2914:420 2914:1206 2914:2203 2914:3200
- BGP Extended Communities Attribute ([RFC 4360](#), February 2006)
 - Adds label, value and longer range displayed as “8/16-bit type:48/56-bit value”
 - Useful for L3VPNs, fewer implementations available
 - Slow adoption rate
 - Cannot see the forest for the trees ([RFC 7153](#))

What Network Operators Use

BGP customer communities

Customers wanting to alter local preference on their routes.

NTT Communications BGP customers may choose to affect our local preference on their routes by marking their routes with the following communities. Our regions are listed [here](#).

Community	Local-pref	Description
(default)	120	customer
65520:nnnn50		only within country <nnnn> (see country list below)
65530:nnnn50		only within region <nnnn> (see region list below)
2914:435	50	only beyond the connected country
2914:436	50	only beyond the connected region
2914:450	96	customer fallback
2914:460	98	peer backup
2914:470	100	peer
2914:480	110	customer backup
2914:490	120	customer default
2914:666		blackhole

RFC 1997
Communities:
Widely Deployed
for Internet
Routing

- RFC 1997 style communities, as they have been used for the past 20 years
- Widely documented in training material, operations procedures, policy documentation
- Required in RFPs and documented in contracts

Below you will find a number of network providers community guides. They are intended for **CUSTOMER** use only. If your network, or your upstream's network, is not a customer of one of these networks, you will NOT be able to use the communities outlined in these guides.

We will make every effort to ensure these guides are up to date, but if there is an update, or a network guide, not reflected here please tell us about it by emailing us at [bgp-guide \(at\) onestep.net](mailto:bgp-guide@onestep.net).

Network AS	Network Name
AS 174	Cogent Communications
AS 209	Qwest Communications
AS 513	CERN - European Organization for Nuclear Research
AS 577	Bell Canada
AS 701	MCI Internet Services
AS 1239	Sprint Business Internet
AS 1270	UUNet DE
AS 1273	Cable & Wireless Plc.
AS 1290	PSINet UK
AS 1299	TeliaSonera International
AS 1759	Sonera
AS 2683	Radio-MSU
AS 2764	AAPT/Connect.com.au
AS 2828	XO Communications
AS 2914	NTT Communications
AS 3212	Triera Internet
AS 3216	Golden Telecom
AS 3239	SURNet - Russia
AS 3257	Tiscali International Network
AS 3292	TDC A/S

Sources: <https://www.us.ntt.net/support/policy/routing.cfm> (AS 2914), <https://onestep.net/communities/>

2017-03-23

CEE Peering Days 2017, Ljubljana

Along Came a Problem

- We knew we'd run out of 16-bit ASNs eventually
- 32-bit ASN work started in January 2001
 - RFC 4893 in May 2007
 - RFC 6793 in December 2012
- RIRs started allocating 32-bit ASNs by request in 2007
- No distinction between 16-bit and 32-bit ASNs now
 - Widely used as edge and transit ASNs
- However, you can't fit a 32-bit value into a 16-bit field
 - Can't use native 32-bit ASNs at all
 - 32-bit ASN owners use private ASNs in communities or some other kludge
 - Creates namespace collisions between ASNs



32-bit ASNs in a 16-bit Field

The Solution

Internet Engineering Task Force (IETF)
Request for Comments: 8092
Category: Standards Track
ISSN: 2070-1721

J. Heitz, Ed.
Cisco
J. Snijders, Ed.
NTT
K. Patel
Arrcus
I. Bagdonas
Equinix
N. Hilliard
INEX
February, 2017

BGP Large Communities Attribute

Abstract

This document describes the Large BGP Communities attribute, an extension to BGP-4. This attribute provides a mechanism to signal opaque information within separate namespaces to aid in routing management. The attribute is suitable for use with all Autonomous System Numbers including four-octet Autonomous System Numbers.

Related Work for 32-Bit ASNs in Communities

- 4-Octet AS Specific BGP Extended Community ([RFC 5668](#), October 2009)
 - RFC 4360 style extended community for 32-bit ASNs displayed as “32-bit ASN:16-bit value”
 - Perceived as a micro optimization
- Flexible BGP Communities ([draft-lange-flexible-bgp-communities](#))
 - December 2002 – August 2010
 - BGP peer community grouping, 32-bit ASNs, plus other stuff
 - No consensus or implementations
- BGP Community Container Attribute ([draft-ietf-idr-wide-bgp-communities](#))
 - July 2010 – March 2017, was Wide BGP Communities Attribute
 - Complementary and comprehensive solution
 - Generalized BGP peer community grouping, 32-bit ASNs, plus other stuff
 - No consensus or implementations, needs time to develop
- No Internet routing communities solution for almost 10 years



Why should I care what color the bike shed is?

IETF Support for BGP Large Communities

- Overwhelming interest on the IDR mailing list
 - Network operators
 - Implementers
- Hundreds of messages and counting on the Working Group adoption thread



Like RFC 1997 Communities, but Larger



Design Goals

- Simply “larger”, that’s it...
 - No added complexity or functionality
 - Extend RFC 1997 communities for 32-bit ASNs
 - Signal an action without losing information about either the origin or the target
- Broadly deployable solution that is available quickly
 - Transitive
- Flexibility for network operators to define their own communities
 - Opaque, may be ignored
- A unique namespace for all 16-bit and 32-bit ASNs
 - Parity and fairness as everyone now can use their globally unique ASN
 - No namespace collisions between ASNs
- Easy to implement
- Easy to adopt
- Easy to remember and tell each other on the phone
 - Canonical representation
 - Especially in an international community with many different languages

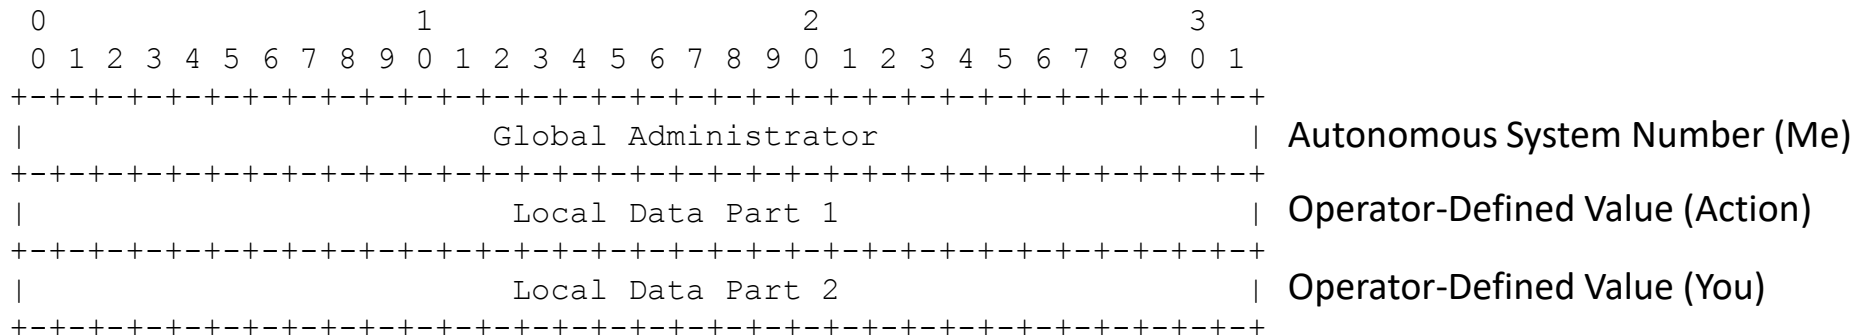
Things That are “Out of Scope”

- No RFC 1997 to BGP Large Communities mapping
 - Out of scope because routing policies differ widely between network operators
- No TLV or header
 - Just use BGP Path Attributes code 32 (0x20)
 - Purposely kept simple to meet the specific use requirements
- No well-known communities
 - Not needed, since RFC 1997 well-known communities like “no-advertise”, “no-export”, “blackhole”, etc. can still be used



*So what'cha
what'cha
what'cha want
what'cha want*

Encoding and Usage



- BGP Large Communities are encoded as a 96-bit quantity and displayed as “32-bit ASN:32-bit value:32-bit value”
- Canonical representation is \$Me:\$Action:\$You
- Working on an [RFC 1998](#) style [draft-ietf-grow-large-communities-usage](#) I-D with examples

BGP Large Community Examples

RFC 1997 (Current)	BGP Large Communities	Action
65400:peer-as	2914:65400:peer-as	Do not Advertise to <i>peer-as</i> in North America (NTT)
43760:peer-as	43760:1:peer-as	Announce a prefix to a certain peer (INEX)
0:43760	43760:0:peer-as	Prevent announcement of a prefix to a certain peer (INEX)
65520:nnn	2914:65520:nnn	Lower Local Preference in Country <i>nnn</i> (NTT)
2914:410	2914:400:10	Route Received From a Peering Partner (NTT)
2914:420	2914:400:20	Route Received From a Customer (NTT)

- No namespace collisions or use of reserved ASNs
- Enables operators to use 32-bit ASNs in \$Me and \$You values

Major Milestones Towards an RFC Standard

Date	Milestone
September 2, 2016	Published draft-heitz-idr-large-community-03
September 6, 2016	Requested IDR WG Adoption
September 24, 2016	IDR Working Group Adoption of draft-ietf-idr-large-community-00
September 29, 2016	Early IANA BGP Path Attributes Code (30) Allocation
October 11, 2016	BGP Large Communities Beacon Prefixes Announced
October 17, 2016	Start of IDR Working Group Last Call
October 26, 2016	Early IANA BGP Path Attributes Code (32) Allocation
November 2, 2016	Start of IETF Last Call and IESG Review
December 1, 2016	Start of IESG Last Call
December 18, 2016	IESG Ballot Issued
January 5, 2017	IESG Approved Revision -12 for RFC Publication
February 16, 2017	RFC 8092 “BGP Large Communities Attribute” Published

RFC 8092

“BGP Large Communities Attribute” Published!

- Idea progressed rapidly from inception in March 2016
- First I-D in September 2016 to RFC publication on February 16, 2017 in just seven months
- Final standard, plus a number of implementation and tools developed as well
- Network operators can test and deploy the new technology now



Timeline Overview



BGP Speaker Implementation Status

Implementation	Software	Status	Details
Arista	EOS	Planned	Feature Requested BUG169446
Cisco	IOS XR	✓ Done!	Beta (perhaps in 6.3.2 for real?)
cz.nic	BIRD	✓ Done!	BIRD 1.6.3 (commit)
ExaBGP	ExaBGP	✓ Done!	PR482
FreeRangeRouting	frr	✓ Done!	Issue 46 (Commit)
nop.nu	freeRouter	✓ Done!	
Juniper	Junos OS	Planned	Second Half 2017 (perhaps 17.3R1?)
MikroTik	RouterOS	Won't Implement Until RFC	Feature Requested 2016090522001073
Nokia	SR OS	Planned	Third Quarter 2017
OpenBSD	OpenBGPD	✓ Done!	OpenBSD 6.1 (commit)
OSRG	GoBGP	✓ Done!	PR1094
rtbrick	Fullstack	✓ Done!	FullStack 17.1
Quagga	Quagga	✓ Done!	Quagga 1.2.0 875
Ubiquiti Networks	EdgeOS	Planned	Internal Enhancement Requested
VyOS	VyOS	Requested	Feature Requested T143

Visit <http://largebgpcommunities.net/implementations/> for the Latest Status

Tools and Ecosystem Implementation Status

Implementation	Software	Status	Details
DE-CIX	pbgpp	✓ Done!	PR16
FreeBSD	tcpdump	✓ Done!	PR213423
Marco d'Itri	zebra-dump-parser	✓ Done!	PR3
OpenBSD	tcpdump	✓ Done!	OpenBSD 6.1 (patch)
pmacct.net	pmacct	✓ Done!	PR61
RIPE NCC	bgpdump	✓ Done!	Issue 41 (commit)
tcpdump.org	tcpdump	✓ Done!	PR543 (commit)
Yoshiyuki Yamauchi	mrtparse	✓ Done!	PR13
Wireshark	Dissector	✓ Done!	18172 (patch)

Visit <http://largebgpcommunities.net/implementations/> for the Latest Status

BGP Large Communities Beacon Prefixes

- The following prefixes are announced with AS path 2914_15562\$
 - 192.147.168.0/24 ([looking glass](#))
 - 2001:67c:208c::/48 ([looking glass](#))
 - BGP Large Community: 15562:1:1

Cisco IOS Output (Without BGP Large Communities Support)

```
route-views>show ip bgp 192.147.168.0
BGP routing table entry for 192.147.168.0/24, version 98399100
Paths: (39 available, best #30, table default)
  Not advertised to any peer
  Refresh Epoch 1
  701 2914 15562
    137.39.3.55 from 137.39.3.55 (137.39.3.55)
      Origin IGP, localpref 100, valid, external
      unknown transitive attribute: flag 0xE0 type 0x20 length 0xC
      value 0000 3CCA 0000 0001 0000 0001
      rx pathid: 0, tx pathid: 0
```

BIRD Output (With BGP Large Communities Support)

```
COLOCLUE1 11:06:17 from 94.142.247.3] (100/-) [AS15562i]
Type: BGP unicast univ
BGP.origin: IGP
BGP.as_path: 8283 2914 15562
BGP.next_hop: 94.142.247.3
BGP.med: 0
BGP.local_pref: 100
BGP.community: (2914,410) (2914,1206) (2914,2203) (8283,1)
BGP.large_community: (15562, 1, 1)
```

BGP Implementer To Do List

- Add support for BGP Path Attributes code 32 (0x20) to BGP
 - Optional CLI command to enable
- Extend your routing policies
 - Set and match
 - Regular expressions
- Extend your show commands
 - Including the debug commands and packet dump output
- Update your documentation
- Update your training material
- Educate your technical staff

Network Operator To Do List


- The entire network ecosystem needs to support BGP Large Communities in order to provision, deploy and troubleshoot
- Ask your routing vendors and implementers for software support
- Update your tools and provisioning software
- Extend your routing policies, and openly publish this information
- Train your technical staff

Image sources: <https://www.sunet.se/blogg/all-i-want-for-christmas-is-large-bgp-communities/>

"All i want for christmas is ... Large BGP Communities" by Fredrik "Hugge" Korsbäck

2017-03-23

CEE Peering Days 2017, Ljubljana



Questions?

Presentation created by:



Greg Hankins

Nokia

greg.hankins@nokia.com

[@greg_hankins](#)



Job Snijders

NTT Communications

job@ntt.net

[@JobSnijders](#)

Visit <http://LargeBGPCommunities.net/> for the Latest Info

Reuse of this slide deck is permitted and encouraged!