Large BGP Communities

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, longer range
  - 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:nnnn	50	only within country <nnnn> (see country list below)
65530:nnnn	50	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

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.

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 Communictions
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

- 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

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

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

IDR
Internet-Draft
Intended status: Standards Track
Expires: April 19, 2017

J. Heitz, Ed.
Cisco
J. Snijders, Ed.
NTT
K. Patel
Arrcus
I. Bagdonas
Equinix
A. Simpson

Nokia
N. Hilliard
INEX
October 16, 2016

Large BGP Communities
draft-ietf-idr-large-community-03

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 in 4-octet ASNs.

Source: https://tools.ietf.org/html/draft-ietf-idr-large-community

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
  - 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
- Wide BGP Communities Attribute (draft-ietf-idr-wide-bgp-communities)
  - July 2010 – September 2016
  - 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?

Image source: http://bikeshed.com/

IETF Support for Large BGP Communities

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

Working Group adoption thread: https://mailarchive.ietf.org/arch/search/?email_list=idr&gbt=1&index=vEa3744YRl5Sj8bUB_I54Uay-fE

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 Large BGP Communities mapping
  - Out of scope because routing policies differ widely between network operators
- No TLV or header
  - Just use BGP Path Attributes code 30 (0x1E)
  - Purposely kept simple to meet the specific use requirements
- No well-known communities
  - Not needed, since RFC 1997 well-known communities like “no-advertize”, “no–export”, “blackhole”, etc. can still be used

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

Encoding and Usage

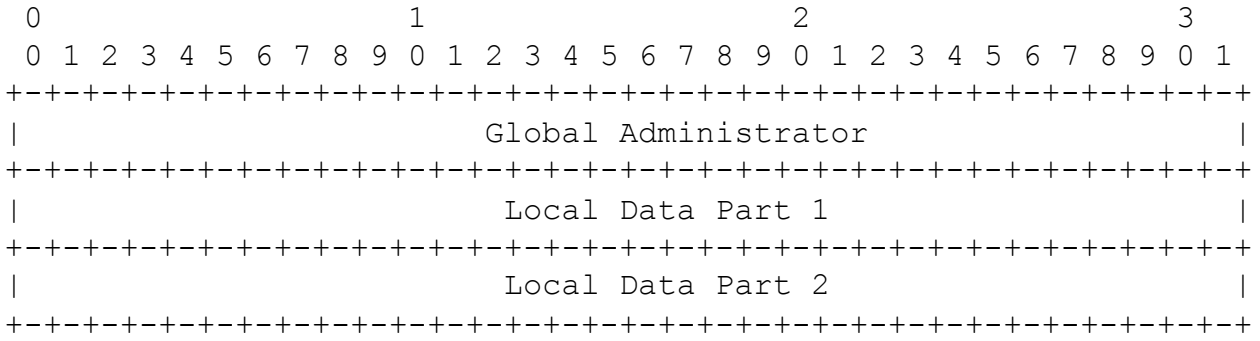

- Large BGP Communities are encoded as a 96-bit quantity and displayed as “32-bit ASN:32-bit value:32-bit value”
- Canonical representation is $Me:$Action:$You
- Working on an RFC 1998 style -usage draft with examples

Large BGP Community Examples

RFC 1997 (Current)	Large BGP Communities	Action
65400:*peer-as*	2914:65400:*peer-as*	Do not Advertise to *peer-as* in North America (NTT)
0:*peer-as*	6667:0:*peer-as*	Do not Announce to Route Server *peer-as* (AMS-IX)
65520:*nnn*	2914:65520:*nnn*	Lower Local Preference in Country *nnn* (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 us 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 1, 2016	Published draft-ietf-idr-large-community-01
October 8, 2016	Published draft-ietf-idr-large-community-02
October 11, 2016	Large BGP Communities Beacon Prefixes Announced
October 16, 2016	Published draft-ietf-idr-large-community-03
October 17, 2016	Start of IDR Working Group Last Call

Timeline Overview

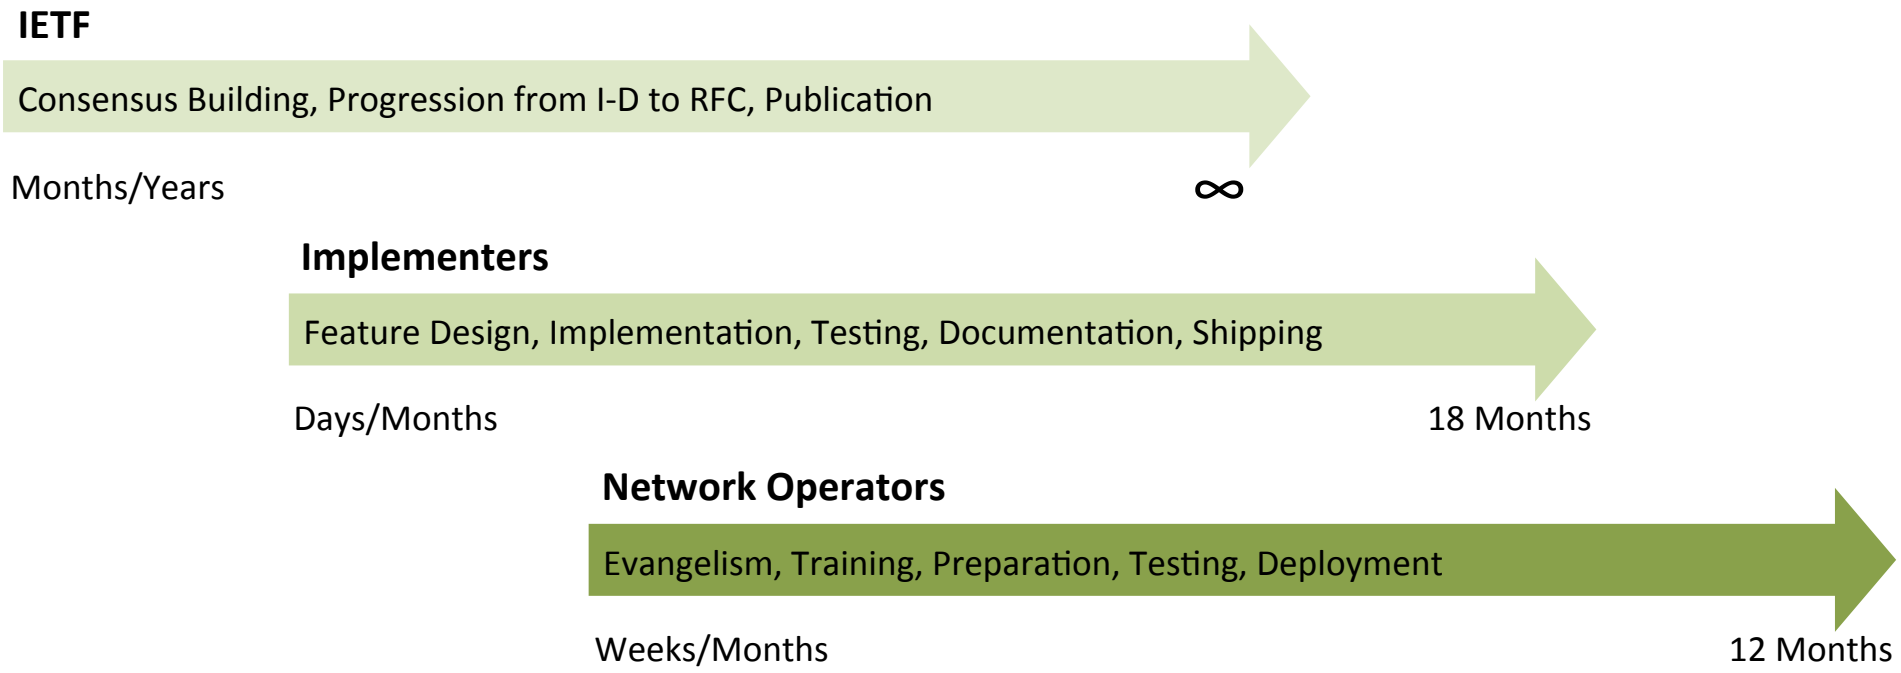

BGP Speaker Implementation Status

Implementation	Software	Status	Details
Arista	EOS	Planned	Feature Requested BUG169446
Cisco	IOS XR	✔ Done!	Engineering Release
cz.nic	BIRD	✔ Done!	BIRD 1.6.3 (commit)
ExaBGP	ExaBGP	✔ Done!	PR482
MikroTik	RouterOS	Won't Implement Until RFC	Feature Requested 2016090522001073
Nokia	SR OS	Planned	
OpenBSD	OpenBGPD	✔ Done!	OpenBSD 6.1 (commit)
OSRG	GoBGP	✔ Done!	PR1094
rtbrick	Fullstack	Planned	ETA: December 2016
Quagga	Quagga	Requested	Feature Requested 875
VyOS	VyOS	Requested	Feature Requested T143

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

Tools and Ecosystem Implementation Status

Implementation	Software	Status	Details
FreeBSD	tcpdump	✔ **Done!**	PR213423
pmacct.net	pmacct	✔ **Done!**	PR61
OpenBSD	tcpdump	✔ **Done!**	OpenBSD 6.1 (patch)
tcpdump.org	tcpdump	✔ **Done!**	PR543 (commit)
Wireshark	Dissector	✔ **Done!**	18172 (patch)

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

Large BGP 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)
  - Large BGP Community: 15562:1:1

Cisco IOS Output (Without Large BGP Communities Support)

```
route-views>sh 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 0x1E length 0xC
        value 0000 3CCA 0000 0001 0000 0001
      rx pathid: 0, tx pathid: 0
```

BIRD Output (With Large BGP 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 30 (0x1E) 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 Large BGP 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

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!