Not so ~~Large~~ BGP Communities

David Freedman

david.freedman@uk.clara.net

Claranet

In brief.

- BGP Communities Attribute (RFC 1997, Aug 1996)
  - Designed for Internet
  - Broad support in BGP implementations.
  - 32 bits, commonly expressed as "**16-bit ASN** : 16-bit VALUE"
  - No space for ASN32 ☹

- BGP *Extended* Communities Attribute (RFC 4360, Feb 2006)
  - Designed for L2/3VPN
  - Limited support in BGP implementations, slow adoption in newcomers.
  - 64 bits, commonly expressed as "8/16-bit TYPE : 48/56-bit VALUE"
  - The VALUE further expressed in subsequent extensions
    - RFC5668 (Oct 2009) specifies an ASN32 TYPE
      - subsequent value of "**32-bit ASN** : 16-bit VALUE" – Not ideal ☹

So what other options for ASN32 users?

- F^{l}e$_{x}$i^{b}l$_{e}$ Communities Attribute (draft-lange-flexible-bgp-communities, Dec 2002)
  - Complex expression, essentially “16-bit TYPE: **32-bit ASN** : up-to-2040-bit VALUE”
  - No consensus or implementations.
- W i d e Communities Attribute (draft-ietf-idr-wide-bgp-communities, Jul 2010)
  - Extremely complex expression, “kitchen sink” approach, 16 bit LENGTH -> up to 65KB VALUE!
  - No consensus or implementations.

- Large Communities Attribute (draft-ietf-idr-large-community, Sep 2016)
  - Simple (by comparison) expression, “**32-bit ASN** : 32-bit VALUE1 : 32-bit VALUE2”
  - **Consensus**, (from both operators, and the IETF), **implementations** (both commercial and open source), and **accepted to be published as an RFC** (currently awaiting a number).

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)

- Canonical representation is *$ME: $ACTION: $YOU*
- No namespace collisions or use of reserved ASNs
- Enables operators to use 32-bit ASNs in *$ME* and *$YOU* values

BGP Speaker Implementation Status

Implementation	Software	Status	Details
Arista	EOS	Planned	Feature Requested BUG169446
Cisco	IOS XR	Planned	Engineering Release, rumour has Second Half 2017
cz.nic	BIRD	✔ Done!	BIRD 1.6.3 (commit)
ExaBGP	ExaBGP	✔ Done!	PR482
Juniper	Junos OS	Planned	Second Half 2017
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	Planned	February 2017
Quagga	Quagga	✔ Done!	Patch Provided for 1.1.0 875
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
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

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 0x20 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)
```

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.
  - **Cisco:** waiting for IOS-XR, but they have no plan for IOS(XE) and NX-OS.
  - **Juniper** have started already, but only Junos OS is planned.
  - **Brocade and Huawei**, no commitments, why??
- **Update** your tools and provisioning software
- **Extend** your routing policies, and openly publish this information
- **Train** your technical staff

Questions?

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

Further questions -> myself or job@ntt.net