



nic.br

Brazilian Network
Information Center

egi.br

Brazilian Internet
Steering Committee



registro.br cert.br cetic.br ceptro.br ceweb.br ix.br

CGI.br members and former members
(only the current members have right to vote)

GENERAL ASSEMBLY

7 members elected by the General Assembly

ADMINISTRATIVE
COUNCIL

AUDIT
COMMITTEE

ADMINISTRATION
.....
LEGAL
.....
COMMUNICATION
.....
ADVISORIES:
CGI.br and PRESIDENT

EXECUTIVE
BOARD

1 2 3 4 5

registro.br

Domain Registration
IP Assignment

cert.br

Security and
Incident Response

cetic.br

Studies and Surveys
About ICT use

ceptro.br

Internet Engineering
and New Projects

ceweb.br

Web Technologies

ix.br

Traffic Exchange

W3C
BRASIL

Web Standards

- 1 Chief Executive Officer
- 2 Administrative and Financial Director
- 3 IT and Services Director
- 4 Director of Special Projects and Development
- 5 Consulting Director for CGI.br activities



1 2 3 4 5 6 7 8 9

GOVERNMENT (Appointed)

10 11 12 13 14 15 16 17 18 19 20 21

CIVIL SOCIETY (Elected)

and

Government Representatives:

- 1 Ministry of Science and Technology (Coordination)
- 2 Presidential Cabinet
- 3 Ministry of Communications
- 4 Ministry of Defense
- 5 Ministry of Development, Industry and Foreign Trade
- 6 Ministry of Planning, Budget and Management
- 7 National Telecommunications Agency
- 8 National Council of Scientific and Technological Development
- 9 National Forum of Estate Science and Technology Secretaries

Civil Society Representatives:

- 10 Internet Expert
- 11 General Business Sector Users
- 12 Internet Service Providers
- 13 Telecommunication Infrastructure Providers
- 14 Hardware and Software Industries
- 15 a 18 Non-governmental Entity
- 19 a 21 Academia

nic.br egi.br

ix.br

BGP Large Communities: a new BGP Routing Heuristic
Montevideo, Uy | 21/09/2017

BGP Large Communities: a new BGP Routing Heuristic

Lucenildo Aquino Júnior

ix.br nic.br cgi.br

Agenda

- BGP Communities
- BGP Extended Communities
- BGP Large Communities
- Comparison between standards
 - BGP Community
 - BGP Extended
 - BGP Large Communities
- References

BGP Communities

- BGP Attribute
 - Type code: 8
 - 32 bits separated in 2 octets.
- Mechanism for labeling route in BGP Protocol
- Well-known communities defined and reserved by IANA.

BGP Communities

- Well-known communities
 - NO_EXPORT - 65535:65281 (0xFFFFFFFF01)
 - Do not export for EBGP neighbors
 - NO_ADVERTISE - 65535:65282 (0xFFFFFFFF02)
 - Do not export for BGP neighbors (EBGP and IBGP)
 - NO_EXPORT_SUBCONFED - 65535:65283 (0xFFFFFFFF03)
 - Do not export for EBGP neighbors, including same confederation

BGP Communities

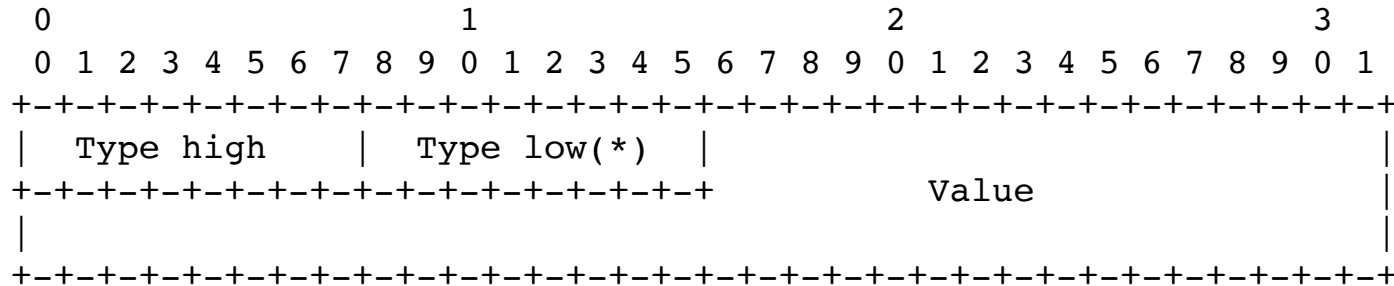
- Well-known communities
 - NO_PEER - 65535:65284 (0xFFFFFFFF04)
 - Do not export learning routes for any private peer
 - BLACKHOLE - 65535:666 (0xFFFF029A)
 - Drop all traffic for defined prefix

BGP - Extended Communities Overview

- Atributo BGP
 - Type code: 16
 - 64 bits separated in 8 octets
- Enhanced mechanism for labeling route in BGP Protocol
- Expand brands range for multiple use cases

BGP - Extended Communities Overview

- Additional a typing field
 - Size: 1 or 2 Octets
 - Enables application-based policer creation



Source: <https://tools.ietf.org/html/rfc4360>

BGP - Extended Communities

Extended Communities Types

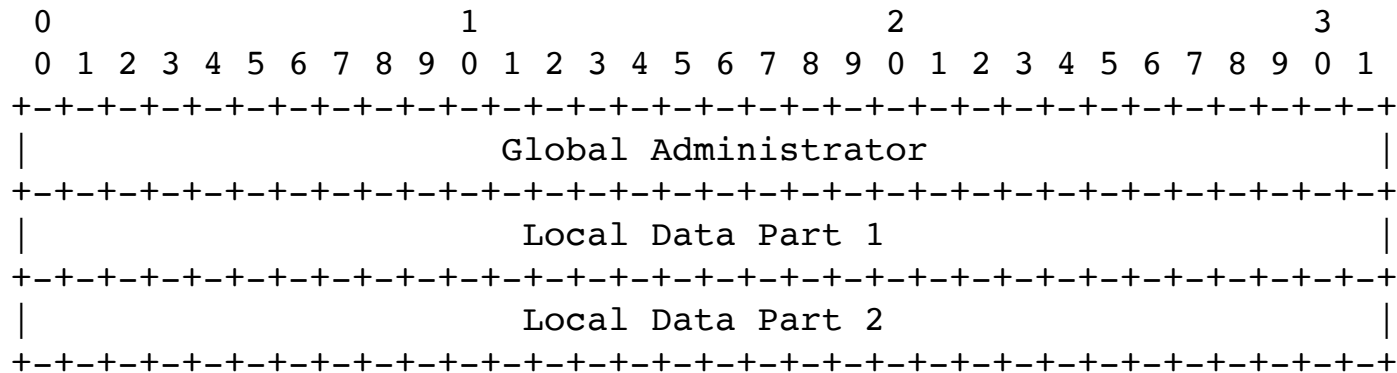
- BGP Extended Communities orientadas a ASNs
 - <TYPE><SUBTYPE>:<GLOBAL_ADM>:<LOCAL_ADM>
 - 16bits:16bits:32bits
- BGP Extended Communities orientadas por IPv4
 - <TYPE><SUBTYPE>:<GLOBAL_ADM>:<LOCAL_ADM>
 - 16bits:32bits:16bits
- BGP Extended Communities Opacas
 - <TYPE><SUBTYPE>:<VALUE>
 - 16bits:48bits

BGP - Large Communities Overview

- Atributo BGP
 - Type code: 32
 - 96 bits separated in 12 octets
 - Defined in Feb. 2017 through RFC 8092
 - Removing type for simplifies use
 - Expand Possibilities
 - \$ASN:\$Function:\$Parameter
 - \$Me:\$Action:\$You

BGP - Large Communities Overview

- Standard BGP communities can be encapsulated
- 2^{96} = Never overlap



Source: <https://tools.ietf.org/html/rfc8092>

BGP - Large Communities

Overview

- In last June was published RFC 8195, that document defined recommendations for large communities use
 - Published by GROW Group for informational purpose
- Large Communities Categorization
 - Informational Communities
 - Action Communities

BGP - Large Communities

RFC - 8195 - Informativos

- Informational Communities

ISO 3166-1 Country Code			ONU M.49			IX.br Location	
Large Community	Desc		Large Community	Desc		Large Community	Desc
26162:1:076	Brazil		26162:2:029	Caribe		26162:20:011	Sao Paulo
26162:1:553	Aruba	+	26162:2:013	Central America	+	26162:20:021	Rio de Janeiro
26162:1:188	Costa Rica		26162:2:005	South America		26162:20:061	Brasilia
26162:1:858	Uruguay		26162:2:019	America		26162:20:666	Non Participant

- For ASN 22548 - NIC.br - "26162:1:076 26162:2:005 26162:20:011"
- For ASN 28000 - Lacnic - "26162:1:858 26162:2:005 26162:20:666"

BGP - Large Communities

RFC - 8195 - Informacionais

- Georeferenced Communities
 - Based in Address

CEP Brazilian Postal Code	
Large Community	Desc
26162:21:05406000	Rua: Teodoro Sampaio, Pinheiros, São Paulo, SP
26162:21:04578000	Av. Nações Unidas, Brooklin Paulista, São Paulo, SP
26162:21:04723003	Av. João Dias, Santo Amaro, São Paulo, SP
26162:21:70070914	SAUS, qd.5, It.6, Asa Sul, Brasília, DF

BGP - Large Communities

RFC - 8195 - Informativos

- Georeferenced Communities
 - Based in Geographic coordinates

GPS Coordinates	
Large Community	Desc
22548:023601489:046694702	NIC.br - NU
22548:023623278:046730221	NIC.br - JD
52462:112096799:068909854	AMX-IX Caribbean
14650:137495038:1127120265	f.dns.br - Korean, Seoul

BGP - Large Communities

RFC - 8195 - Action

- ASN Based Selective No-Export

No-Export	
Large Community	Desc
26162:4:22548	Não exporta rotas para o AS22548
26162:4:14026	Não exporta rotas para o AS22548
26162:4:11752	AMX-IX Caribbean

BGP - Large Communities

RFC - 8195 - Action

- Location Based Selective No-Export

Selective No-Export	
Large Community	Desc
26162:5:076	Não exporta rotas para EBGP localizados no Brasil
26162:5:553	Não exporta rotas para EBGP localizados em Aruba
26162:5:858	Não exporta rotas para EBGP localizados no Uruguai

BGP - Large Communities

Deploy Status

- Vendors

Vendor	Software	Status	Details
Arista	EOS	Planned	Feature Requested BUG169446
Cisco	IOS XE	Planned	16.9.1 (FCS July 2018) (source)
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)
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
ndp.hu	freeRouter	✓ Done!	
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	EdgeOS	Planned	Internal Enhancement Requested
VyOS	VyOS	Requested	Feature Requested T143

Source: <http://largebgpcommunities.net/implementations/>

BGP - Large Communities

Deploy Status

- Ecosystems and Tools

Vendor	Software	Status	Details
DE-CIX	phgpp	✓ 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)

Source: <http://largebgpcommunities.net/implementations/>

BGP - Large Communities Deploy Status

The image shows a Wireshark capture of a BGP UPDATE message. The packet list pane shows a single packet (No. 1) at time 07:21:26.912456, source 192.0.2.2, destination 192.0.2.3, protocol BGP, length 441 bytes. The packet details pane shows the following structure:

- Next hop: 192.0.2.2
- Path Attribute - LARGE_COMMUNITY: 65535:1:1 4294967295:4294957295:4294967295
 - Flags: (xrd), optional, transitive: optional, transitive, complete
 - Type Code: LARGE_COMMUNITY [32]
 - Length: 24
 - Large Communities: 65535:1:1
 - Global Administrator: 65535
 - Local Data Part 1: 1
 - Local Data Part 2: 1
 - Large Communities: 4294967295:4294967295:4294957295
 - Global Administrator: 4294967295
 - Local Data Part 1: 4294967295
 - Local Data Part 2: 4294967295
- Network Layer Reachability Information (NRI)

The packet bytes pane shows the raw data in hexadecimal and ASCII. The ASCII column shows the following characters: .B...B...E... .960.E. | |... ..M. 3... ..?/... ..K... /0... ..N...

BGP Communities Comparison

	BGP Standard Community	BGP Extended Community	BGP Large Community
ASN 16bits			
ASN 32bits			
ASN 32bits in Both Sides			
Sobreposição			

References

- Chandra, R., Traina, P., and T. Li, "BGP Communities Attribute", [RFC 1997](#), August 1996.
- Rekhter, Y. and T. Li, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.
- Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4364](#), February 2006.
- IANA, "BGP Path Attributes", <<http://www.iana.org/assignments/bgp-parameters>>.
- Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", [RFC 4360](#), February 2006.
- Heasley, J., Simpson, A., "BGP Large Communities Attribute", [RFC 8092](#), February 2017
- Snijders, J., Heasley, J., Schmidt, M., "Use of BGP Large Communities", [RFC 8195](#), June 2017
- Large Communities Initiative, <<http://largebgpcommunities.net/>>

Thanks!

<http://ix.br>



lucenildo@nic.br



[@lucenildolaj](https://twitter.com/lucenildolaj)



<http://fb.com/nic.br>

21 de Setembro de 2017

nic.br **cgi.br**

www.nic.br | www.cgi.br