

# BGP Communities for service providers

Marco d'Itri

<md@seeweb.it>

@rfc1036

Seeweb s.r.l.

Albanian Network Operators Group meeting - 14 November 2018



# Attributes of BGP routes

## Mandatory attributes

- as-path
- next-hop
- origin

## Some optional attributes

- Community: a 32 bit (or more...) integer, can have multiple values.
- Local preference: a 32 bit integer.
- Multi-exit discriminator (MED): a 32 bit integer.

# How routers choose the best path

In this order:

- Highest local preference.
- Shortest as-path.
- Lowest origin<sup>1</sup>.
- Lowest MED.

(And some other parameters which are less important and may be implementation-dependent.)

---

<sup>1</sup> It should be rewritten as egp on ingress anyway!

# MED and local preference

## Multi-exit discriminator

If received over eBGP then it will not be sent to other eBGP neighbors.  
Often accepted only from customers.  
May be reset or changed on ingress for traffic engineering purposes.

## Local preference

Never sent to eBGP neighbors.  
Used to create your local routing policy.

# Types of communities

Regular: 32 bit, represented as 16+16 bit.

Extended: 64 bit.

Large: 96 bit, represented as 32+32+32 bit.

Nobody cares about extended communities, so they will not be discussed further here.

# Representation of communities

## Regular: 16+16 bit

- 12637:1234
- 0:12637
- 1:2

## Large: 32+32+32 bit

- 43369:1:250000
- 1:2:3
- 250000:1:2

Usually administrators encode their own ASN in the first pseudo-field for communities intended to be propagated to other ASes.  
32 bit ASNs can be encoded only in large communities!

# How are communities used?

Some examples:

Received from neighbors (or your other routers)

Influence your routing policy.

Advertised to neighbors

Influence the routing policy of your transit providers.

# Tagging communities (inbound)

Route received, e.g.:

- from a customer.
- from a peer.
- in a specific city/country/region.

This kind of community can be evaluated in your own routing policy, e.g. to set an appropriate localpref.

# Traffic engineering (outbound)

Generally transit providers provide TE communities to their customers, but peers do not.

## Example actions:

- Prepend N times
- Do not announce
- Change the local preference
- Blackhole all traffic to the prefix

## Actions can be limited, e.g.:

- To a specific transit provider or peer of your transit provider.
- To a specific IXP.
- In a specific country or region.

## By IXPs route servers

IXPs often implement these communities in their route servers to allow customers to control how the RSEs should readvertise the routes.

If the RS ASN is 43369, then:

- `0:PEERASN` or `43369:0:PEERASN`: do not advertise to PEERASN.
- `43369:PEERASN` or `43369:1:PEERASN`: advertise to PEERASN.
- `0:43369` or `43369:0:0`: do not advertise to any peer.

# Any questions?



<https://www.linux.it/~md/text/communities-alnog2.pdf>  
(Google ... Marco d'Itri ... I feel lucky)

