

Large BGP Community

draft-ietf-idr-large-community-07

IETF 97, November 2016
Seoul

Jakob Heitz (Cisco)
Job Snijders (NTT)
Keyur Patel (Arrcus)
Ignas Bagdonas (Equinix)
Nick Hilliard (INEX)

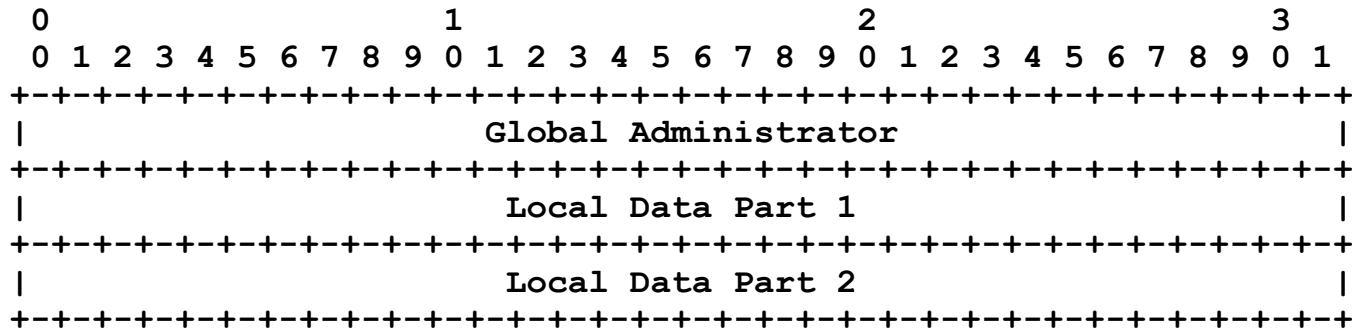
Benefits

- Today's RFC1997 communities are 4 octets each. 2 for ASN and 2 to indicate the action.
- With 4 octets for the ASN, it is mapped one to one to your actual globally unique assigned ASN. Clean name space!
- With 8 octets for the action, you can easily create opaque action code points: you have room for a "target" and an "action"

Example:

- "2914:303:199036" (In NTT: prepend 3 times to 32-bit AS 199036)
 - **(currently not possible** with NTT's standard BGP Communities)

Encoding



- GA is an ASN... usually.
- The owner of the ASN defines the Local Data Part.
- Uses:
 - Others tell the ASN owner how to treat the route.
 - ASN owner tells others properties of the route.
 - Community is used privately: GA can be anything.

Canonical Representation

- Each large community is represented as a:b:c
- a, b and c are each a non-negative decimal integer
 - a – ASN
 - b – Local Data 1
 - c – Local Data 2
- Why it matters
 - ISPs publish their communities in a consistent format.
 - <https://onestep.net/communities/>

The road was rough


~~30~~ → ~~31~~ → 32

”Dogfooding” is **very** important

After early allocation of a **transitive optional attribute**, immediately do the following:

- Inform networking community you are putting up a beacon
- Announce beacon into DFZ
- Look at RIPE Atlas, RIPE RIS, routeviews, etc
- Put real users behind it

The test lab



draft-ietf-idr-large-community-07

Issue with 30

- Huawei implemented a different feature on Attribute-30
- The implementation considered the attribute malformed, because the Large Communities length did not match their expected length
- Treat-as-withdrawal

Conclusion: The Large Communities deployer was punished, thus incentive is for Large to renumber

Further reading: <https://labs.ripe.net/Members/emileaben/large-bgp-communities-largely-accepted-now>

Future work

1. NTT (Jared + Job) will perform controlled probing of the Attribute space in 2017. This will be a longer term, slow paced, well documented & published project.
2. Clean up existing known garbage: **draft-snijders-idr-deprecate-30-31-129** (Which also covers 241, 242, 243)
3. When you get an early allocation, make sure you are in a position to immediately test and if necessary return the codepoint