

# **BGP and its Communities**

Large and Small

# Problem space

- AS4 is being deployed and is visible in AS paths.
- Policy needs to be expressed with AS4 elements.
- Communities are the fundamental building block of policy signalling.
- No simple mechanism to encode AS4:AS4 policy constructs.
- There is a multitude of hacks to work around this problem, although no clean solution.

# Hacks

- Split AS4 into AS2 + AS2 and use two standard communities.
- AS2 portions may conflict with legitimate AS2s – this is too fragile.
- Remap AS4 to unused AS2. It works if you can find an unused AS2.
- Extended communities allow to carry AS4:AS2, in some use cases that might be enough.
- Extended communities have types – all parties need to agree on the type to use.

# Examples of Hacks

- Around 1100 of different instances of extended communities in ipv4u address family in global routing table.

```
extcomm RT:65165:132167 (Route Target) (type 0/2) (Two-Octet AS Specific Extended Community)
extcomm RT:65171:132203 (Route Target) (type 0/2) (Two-Octet AS Specific Extended Community)
# AS4 signalling
extcomm RO:197801:64515 (Route Origin) (type 0/3) (Two-Octet AS Specific Extended Community)
# AS4 signalling but different type
extcomm RT:65204:222220000 (Route Target) (type 0/2) (Two-Octet AS Specific Extended Community)
extcomm RT:65228:222220000 (Route Target) (type 0/2) (Two-Octet AS Specific Extended Community)
# this seems to be some kind of marker
extcomm RT:65300:29025 (Route Target) (type 0/2) (Two-Octet AS Specific Extended Community)
extcomm RT:9255:6010 (Route Target) (type 0/2) (Two-Octet AS Specific Extended Community)
# extended communities used for AS2 signalling
extcomm OSPF-DI:0:131584 (OSPF Domain Identifier) (type 0/5) (Two-Octet AS Specific Extended Community)
extcomm OSPF-RT:000000000500 (OSPF Route Type) (type 3/6) (Opaque Extended Community)
extcomm OSPF-RT:0:1234736768 (OSPF Route Type) (type 0/6) (Two-Octet AS Specific Extended Community)
# some creative uses or possibly filtering problems
extcomm L2VPN-ID:15.15.15.15:0 (L2VPN Identifier) (type 1/10) (IPv4 Address Specific Extended Community)
extcomm L2VPN-ID:16.16.16.16:0 (L2VPN Identifier) (type 1/10) (IPv4 Address Specific Extended Community)
# someone's fabric infrastructure leak?
```

# Large Communities

Hacks are not sustainable – we need something better. That something would be large communities.

- A simple approach continuing along the standard communities.
- Larger fields, more fields, and a clean namespace separation.
- Global and two local parts, all treated as 32 bit integers.
- Policy configuration stays similar to but not the same as standard communities.
- Simple and practical approach.

# Large Communities - encoding

Global Administrator
Local Data Part 1
Local Data Part 2

- New type of BGP path attribute, currently allocated value is 32.
- No capabilities.
- Not directly backwards compatible with standard or extended communities, no mapping defined.
- No Well Known large community values.
- Global Administrator values of 0, 0xffff, and 0xffff`ffff are reserved for extensibility – similar to reserved AS numbers.
- Not a direct replacement of standard communities – this is intentional.

# Large Communities - policy

Policy format is ASN:X:Y, ASN defines a namespace for X and Y interpretation. Canonical representation format is A:B:C, implementations may support other formats as long as it is treated as 3 x 32 bit fields of data.

X and Y can be used in any practical way – there is enough room to encode the source and the target of the policy action. Bidirectional signalling is possible.

29535:201054:123456 – thank you TPIX for hosting the Euro-IX meeting.

- ASN:To Upstream1:Prepend
- ASN:To Upstream2:Do not advertise
- ASN:From Upstream1:Received at edge 11
- ASN:From Upstream2:Received at edge 22

# Large Communities - policy

- Large communities are transitive.
- Propagation may be controlled in the same or different way compared to standard communities – this is an implementation detail.
- Large and standard communities do not need to be propagated together in the same way – this is a local policy detail.
- Adding and removing of large communities is controlled by the routing policy. Large communities have meaning in the namespace ASN and may have meaning in other ASNs, however that is not mandatory.

# Large Communities - examples

6667:0:123456 – AMS-IX do not advertise to peer 123456

3356:123456:666 – LVLT advertise blackhole to AS123456 only.

3356:666:123456 – it is up to AS3356 to interpret the policy in this way with the same meaning.

3356:400:10 2914:3356:65400 2914:400:20 – a prefix was received from LVLT peering partner NTT in North America which received that prefix from its customer.

# Going Forward

- Standard communities will stay for a foreseeable future – at least due to well known community values.
- Extended communities will continue to be used for specific address families, AS4 problem does not go away and extensions there will be needed. It is out of scope of large communities.
- Large communities will be initially deployed for ipv4u/ipv6u address families.
- Large communities proposal does not obsolete wide communities proposal – wide communities address a different set of problems and will coexist with the other types.
- Many vendors have firm plans for large community support – some already do, some will do in a near future. Talk to your favourite vendors and ask for large communities to be implemented.
- There are prefixes in global routing system already carrying large communities, and there are operators able to act on large communities too.

# Large Communities Summary

- It is a larger and more flexible version of standard communities.
- No specific format enforcement, operators are free to define their own routing policies in a flexible and practical way.
- 32-bit ASN clean solution, no conflicts and overlapping values.
- Not complex to implement, not complex to operate.

<http://largebgpcommunities.net>