

Большие сообщества BGP

Large BGP Communities

BGP communities

- Communities represent argument abstraction mechanism for BGP policy propagation.
- Arbitrary dimensionless value, typically having AS-wide namespace.
- Well known and universally deployed.
- ASN2 clean only, no simple way of expressing policy with ASN4 constructs.
- No practical mechanism for expressing AS4 clean routing policy.

Examples of 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 can represent AS4:AS2, in some use cases that might be enough.
- Extended communities have types – all parties need to agree on the type to use.

Around 1200 of different instances of incompatible extended communities in ipv4u address family in global routing table.

```
Extcomm ASN4:RO:197999:0 (Route Origin) (Four Octet AS Specific)
Extcomm ASN4:RT:196844:7777 (Route Target) (Four Octet AS Specific)
Extcomm ASN2:RT:0:201494 (Route Target) (Two Octet AS Specific)
Extcomm ASN2:RO:0:198570 (Route Origin) (Two Octet AS Specific)
Extcomm ASN4:SAS:202867:40001 (Source AS) (Four Octet AS Specific)
Extcomm ASN4:RO:203507:2914 (Route Origin) (Four Octet AS Specific)
```

Flexible communities

Structure	Type	ASN
ASN		Length
Value		

- AS4:AS4 flexible community would be encoded in ASN:Value fields.
- Flexible communities define neighbour class for controlling attribute scope propagation (Peer, Customer, Upstream, Confederation, All).
- It is a generic container for carrying multiple information objects, not just ASNs.
- Proposed around 2002 timeframe.
- No relevant implementations.

Wide communities

Type	Flags	Hop Count
Length	Reg/Loc Type = 1	
Reg/Loc Type = 1	Source AS	
Source AS	Context AS	
Context AS	AType = 4	ALen = 4
AVal		

- AS4:AS4 wide community would be encoded in ContextAS:AVal fields.
- 24 octets on the wire to represent 8 octets of information
- Nontrivial parsing increases implementation complexity
- Propagation scope control is even more elaborate

- Proposed around 2006 timeframe
- No relevant implementations

BGP Large community attribute

Hacks are not sustainable – we need something better.

- A simple approach continuing along the standard communities.
- Larger fields, more fields, and a clean namespace separation.
- Policy configuration stays similar to standard communities.
- Simple and practical approach.
- RFC 8092, standardization process is now complete.
- Some implementations available, many on the way.
- Large communities are already present in the routing system.

Large communities - encoding

Global Administrator
Local Data Part 1
Local Data Part 2

- New type of transitive BGP path attribute.
- No extensions, no capabilities, no options – simple container.
- 3 x 32 bit fields interpreted as integers.
- Not directly backwards compatible with standard or extended communities, no mapping defined – this is intentional.
- No Well Known large community values.
- Global Administrator values of 0, 0xffff, and 0xffff`ffff are reserved for extensibility – similar to reserved AS numbers.

Large communities - policy

- Policy format is ASN:X:Y, ASN defines the 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.
- Common agreement on Context:Function:Argument interpretation.
- Large communities have meaning in the namespace context ASN and may have meaning in other ASNs, however that is not mandatory.

Large communities - policy

- Informational LC – tracking and recording of information. ASN indicates the originator of information.
- Action LC – requesting specific action to be taken. ASN indicates the action target.
- No predefined equivalent of NO_EXPORT/NO_ADVERTISE. This needs to be implemented as part of routing policy.
- Flexible geolocation encoding. ISO 3166-1 numeric country identifiers and UN M.49 region identifiers could be directly mapped to LC local part fields.
- Other location information can be transcoded into numeric format – UK postal codes as an example.

Large Communities - examples

- 8631:0:123456 – MSK-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.

Standard	Large	Function
65400:peer-as	2914:65400:peer-as	NTT: do not advertise to peer-as in North America
43760:peer-as	43760:1:peer-as	INEX: Advertise to peer-as
0:43760	43760:0:peer-as	INEX: Do not advertise to peer-as
65520:nnn	2914:65520:nnn	NTT: Lower Local Preference in Country nnn

Large communities - compatibility

- Large communities do not replace standard communities.
- Large communities are not backwards compatible with standard communities.
- For BGP, standard and large communities are different objects. For a user, standard and large communities are containers of information.
- A function can be defined for large communities that carries the value of standard communities.
- ASN:0:Value might be one possible option.
- SC ASN:X represented by LC ASN:0:X

Going forward

- Standard communities will stay relevant for a foreseeable future.
- Extended communities will continue to be used for specific address families, AS4 problem does not go away and extensions will be needed. It is out of scope of large communities.
- Large communities do not obsolete wide communities – wide communities address a different set of problems.
- Many vendors have firm plans for large communities support – some already do, some will do in the near future.
- There are prefixes in global routing system already carrying large communities, and there are operators able to act on large communities too.

Things to remember

- Large communities need to be supported in the entire routing system.
- Tools and systems need to be made aware of large communities – frequently this is the most complex part.
- Routing policies need to be adapted. Ideally common agreements on functions and interpretations need to be developed.
- Route policies do not have to migrate away from standard communities, they need to be adapted to take benefit of large communities. Policies should be kept compatible where practical, at least before large communities support becomes globally widespread.

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>