BGP, community tagging and route selection

Bruce.Morgan@aarnet.edu.au

aarnet
Australia's Academic and Research Network

Routing Policy

- We have learned that we can alter the route selection process by changing some of the attribute values of BGP routes
  - If the local preference value is increased then we can favour one route over another – even if the AS-PATH length is longer. This will apply to the way that we see inbound routes
  - If the MED is lowered on the routes advertised then we can generally favour a specific inbound traffic path to our networks
  - AS-PATH prepending can be used to favour less a particular route as when local preference is equal then route determination is based on AS path length.

Selective Routing Policy

- So far we have looked at all routes received from a particular peer and applied a policy either preferring them or not
- In the internet we might want to be selective about the routes we apply policy to, even though they may be only certain routes supplied by a few peers.
  - An example of this may be to favour routes originated from Abilene, APAN or TEIN2 as they have better bandwidth available at low cost.

BGP Attributes: Community

- The community attribute is a mechanism for colouring or tagging routes
- The actual community attribute is a 4 byte word
- There are well known attributes such as
  - No-export 0xFFFFFF01
  - No-advertise 0xFFFFFF02
- Generally community tags are split into two 2 byte values
  - AS:VAL
  - This makes them readable!

Community tags

- Communities can be added to a particular prefix using route-maps:
  - Set community 100:10
  - Set community 100:10 additive
- Communities can be matched by route-maps:
  - ip community-list COMMODITY-PEER 10:1
  - Match community COMMODITY-PEER
- Communities can be deleted from particular prefixes
- In short we have control over community tags!

How do we identify a route?

- We can identify routes sourced by Abilene as they will have _11537_ in their path
  - But, anyone can inject a route with that in it!
- Better still, at a direct peering with Abilene we accept routes that begin ^(11537_)
  - At that point we can apply a BGP community attribute to those routes so that we can identify them when we need to.
  - For example if we peer with Abilene as AS4621 we can tag all routes from Abilene with the community
    - 4621:11537

Route information

- To continue we can even enter even more specific information at that point:

```
as-path access-list 1 permit ^(11537_)
route-map ABILENE-IN permit 10
 match as-path 1
 set community 4621:11537, 4621:1001, 4621:2043
```

The community 4621:1001 is used to signify it is an Research route. AS4621 might have the following routes marked by community tags:

4621:1000	customer institutions
4621:1001	overseas research networks
4621:1002	domestic peers
4621:1003	international peers
4621:1004	commodity internet

Route information

- As iBGP uses loopback addresses then we might devote a /24 address space for all loopbacks of our routers. This would allow us to specify at which router we received this routing information by marking it 4621:2xxx:
  - 4621:2043 might mean the router with the last octet of the ip address being 43. We would then know where the route originated on our network

But routes are only our outbound path…

- Even if we do make a route policy selection favouring specific routes, how do the owners of those routes know how to get back to us – will they traverse commodity or research links?
- We need to identify research organisations within our network and mark them with appropriate tags to pass back so that the right decisions can be made.
- Normally research paths will have a shorter AS path – but that might not always be the case…

At our peering edge..

- Match only the research institutions and advertise them onwards to the international research networks:

```
ip community-list standard NATIONAL-RESEARCH 4621:1000

route-map NATIONAL-RESEARCH-OUT permit 10
 match community NATIONAL-RESEARCH
 set metric 10
!
route-map NATIONAL-RESEARCH-OUT deny 20
!
router bgp 4621
 neighbor 1.1.1.1 route-map RESEARCH-OUT out
```

At our peering edge inbound…

```
ip community-list standard ABILENE-PARTICIPANT 11537:950
ip community-list standard ABILENE-ITN 11537:2501
ip as-path access-list 10 permit ^(11537_)

route-map ABILENE-RESEARCH permit 10
match community ABILENE-PARTICIPANT
match as-path 10
set local preference 120
set community 4621:1001
!
route-map ABILENE-RESEARCH permit 20
match community ABILENE-ITN
match as-path 10
set local preference 120
set community 4621:1001
```

And at another peering edge…

```
ip community-list standard TEIN2-PARTICIPANT 29999:30
ip community-list standard TEIN2-TRANSIT 29999:40
ip as-path access-list 10 permit ^(29999_)

route-map TEIN2-RESEARCH permit 10
match community TEIN2-PARTICIPANT
match as-path 10
set local preference 120
set community 4767:1001
!
route-map TEIN2-RESEARCH permit 20
match community TEIN2-TRANSIT
match as-path 10
set local preference 120
set community 4767:1001
```

And at the other side of the peering..

- Normal practice would be to give a higher local preference for routes advertised by another research network
  - This could be open
    - route-map LEARNED_FROM_RESEARCH_PEER permit 10
      - set local preference 120
  - Or could be refined
    - route-map LEARNED_FROM_RESEARCH_PEER permit 10
      - match community 4621:1001
      - match as-path 10
      - set local preference 120

A problem to solve

- We have two international research peers with routes tagged.
- One customer wants only research routes.
- The other customer is not a research institution and should not get research routes from international destinations but can get access to domestic research institutions.