



E Q U I N I X

Routing Control at Peering Points

HKNOG 0.1

Raphael Ho

Peering Points Types

- **Bi-Lateral**
 - Fine grained control
 - Potentially complicated to manage
- **Multi-lateral**
 - Coarse grained control
 - Simple to manage
 - for some definition of ‘manage’
- **Multi-lateral with peering options**
 - Fine grained control
 - Simple to manage

Warning

- Configurations are simplified to illustrate relevant commands
- DO NOT USE AS IS
- You' ll want to add things like
 - Outbound Filters
 - Ensure address aggregation, filter subnets
 - Community Filters, announce customer routes only
 - Bogons, martians and other bad stuff
 - Inbound Filters
 - Max Prefixes
 - iACLs
 - Bogons, martians, specifics and other bad stuff
 - Also turn off things like CDP, STP, VTP, Proxy-ARP, IGP routing announcements, etc for good peering point etiquette

Example: DE-CIX

<http://apps.db.ripe.net/whois/lookup/ripe/aut-num/AS6695.html>

block of announcement of a route to a certain peer 0:peer-as

announcement of a route to a certain peer 6695:peer-as

block of announcement of a route to all peers 0:6695

announcement of a route to all peers 6695:6695

There are additional BGP communities for controlling announcements

- 6695:65000 peer is using bad netmask (80.81.192.0/24)

- 6695:65001 peer is using bad netmask (80.81.193.0/24)

- 6695:65003 prefix is injected at DE-CIX1

- 6695:65004 prefix is injected at DE-CIX2

- 6695:65005 prefix is injected at DE-CIX3

- 6695:65006 prefix is injected at DE-CIX4

- 6695:65281 set community NO-EXPORT

Please take care when setting 0:x communities as this changes the local-preference:

Community	local-preference
-----------	------------------

0:6695	0
0:x, x!= 6695	50
else or nothing	100

Routing Control (1)

announcement of a route to all peers 6695:6695

block of announcement of a route to a certain peer 0:peer-as

- **Default Open Policy**

- Announce to all peer,
except certain routes

```
router bgp 40
  neighbor 80.81.192.1 remote-as 6695
  neighbor 80.81.192.1 route-map decix-in in
  neighbor 80.81.192.1 route-map decix-out
  out
```

- **Example**

- Announce all routes to
MLPE participants,
except to AS10, AS20,
AS30

```
route-map decix-out permit 10
  set community 6695:6695 0:10 0:20 0:30
```

```
route-map decix-in deny 10
  match as-path 10
```

```
ip as-path access-list 10 permit ^10
ip as-path access-list 10 permit ^20
ip as-path access-list 10 permit ^30
```

Routing Control (2)

announcement of a route to a certain peer 6695:peer-as
block of announcement of a route to all peers 0:6695

- **Default Closed Policy**

- Announce only to selected peers

```
router bgp 40
  neighbor 80.81.192.1 remote-as 6695
  neighbor 80.81.192.1 route-map decix-in in
  neighbor 80.81.192.1 route-map decix-out out
```

- **Example**

- Announce only to AS10, AS20, AS30

```
route-map decix-out permit 10
  set community 0:6695 6695:10 6695:20 6695:30
```

```
route-map decix-in permit 10
  match as-path 10
```

```
ip as-path access-list 10 permit ^10
ip as-path access-list 10 permit ^20
ip as-path access-list 10 permit ^30
```

Routing Control (3)

- **Mixed policy**
 - Open for local routes
 - Selective for International routes
- **Example**
 - Announce international only to AS10, AS20, AS30

```
router bgp 40
  neighbor 80.81.192.1 remote-as 6695
  neighbor 80.81.192.1 route-map decix-in in
  neighbor 80.81.192.1 route-map decix-out out
```

```
route-map decix-out permit 10
  match community 40:4000 # local routes
  set community 6695:6695
```

```
route-map decix-out permit 20
  match community 40:5000 # international routes
  set community 0:6695 6695:10 6695:20 6695:30
```

```
route-map decix-in permit 10
  match as-path 10
  set community 40:5000
```

```
route-map decix-in permit 20
  set community 40:4000
```

```
ip as-path access-list 10 permit ^10
ip as-path access-list 10 permit ^20
ip as-path access-list 10 permit ^30
```

Route Preferencing

Communitylocal-preference

0:6695 0

0:x, x!= 6695 50

else or nothing 100

- Useful for load balancing multiple, redundant links
- Example
 - Send Local routes
 - Backup for Intl routes

```
router bgp 40
  neighbor 80.81.192.1 remote-as 6695
  neighbor 80.81.192.1 route-map decix-in in
  neighbor 80.81.192.1 route-map decix-out out
```

```
route-map decix-out permit 10
  match community 40:4000 # local routes
  set community 6695:6695 # default lpref = 100
```

```
route-map decix-out permit 20
  match community 40:5000 # international routes
  set community 0:65000 # local pref = 50
```


Legacy Patch

6695:65000 peer is using bad netmask (80.81.192.0/24)
6695:65001 peer is using bad netmask (80.81.193.0/24)

- DECIX used to run /24 subnet
- Upgraded to /23
- Now at /22
- Informational service to prevent blackholing

```
router bgp 40
  neighbor 80.81.192.1 remote-as 6695
  neighbor 80.81.192.1 route-map decix-in in
  neighbor 80.81.192.1 route-map decix-out out
```

```
route-map decix-in deny 10
  match community 6695:65001
```

```
route-map decix-in permit 20
```

Other information service

6695:65003 prefix is injected at DE-CIX1
6695:65004 prefix is injected at DE-CIX2
6695:65005 prefix is injected at DE-CIX3
6695:65006 prefix is injected at DE-CIX4

- **Distributed Exchange**

```
router bgp 40
  neighbor 80.81.192.1 remote-as 6695
  neighbor 80.81.192.1 route-map decix-in in
  neighbor 80.81.192.1 route-map decix-out out
```
- **Allow multihomed participants to exchange traffic locally without transiting backbone**

```
route-map decix-in permit 10
  match community-list 10
```
- **Example**
 - Accept DE-CIX1 and DE-CIX2 routes

```
ip community-list 10 permit 6695:65003
ip community-list 10 permit 6695:65004
ip community-list 10 deny any
```

Prepending

Prepend TIX AS8235 to *all* TIX-Peers

set community 8235:65500

Prepend TIX AS8235 to *selected* TIX-Peer N-times

set community 6550N:AAAA (N can be 1 to 3)

- Allow granular traffic control across multiple exchanges

```
router bgp 40
  neighbor 194.42.48.40 remote-as 8235
  neighbor 194.42.48.40 route-map tix-in in
  neighbor 194.42.48.40 route-map tix-out out
```

- Example

- Prepend 1 time to everybody
- Prepend 3 times to AS10 only

```
route-map tix-out permit 10
  set community 8235:65500 65503:10
```

Blackhole Peering

set community 24115:666

- Traffic will be null-routed at IX
- Participants may accept 3rd party blackhole routes to reduce backbone load
- Example
 - Announce blackhole for 1.2.3.4
 - Accept blackholes from peers

```
router bgp 40
  neighbor 119.27.63.40 remote-as 24115
  neighbor 119.27.63.40 route-map eqixhk-in in
  neighbor 119.27.63.40 route-map eqixhk-out out
```

```
route-map eqixhk-out permit 10
  match community 40:666 # internal blackhole
  comm
  set community 24115:666
```

```
route-map eqixhk-in permit 10
  match community 24115:666
  set community 40:666
  set ip nexthop 192.168.0.1 # internal null
  route
```

```
interface null0
  ip address 192.168.0.1 255.255.255.255
```

Equinix Internet Exchange Features

<https://ix.equinix.com/ixp/mlpeCommunityInfo>

Definition	Community String
Default Open Policy – Announce to all except to AS12345	24115:24115 0:12345
Default Closed Policy – Announce to none except to AS12345	0:24115 24115:12345
Prepend once to AS12345	65501:12345
Prepend twice to AS12345	65502:12345
Prepend three times to AS12345	65503:12345



EQUINIX

Questions?