flokinet-016 - BGP peering

Nicholas Morrison nick@nanocat.net

Connecting

SSH to the netlab server:

```
$ ssh flokilab.nanocat.net
```

List and connect to the running containerlab devices:

```
$ list-devices
$ connect device-name
```

Diagram

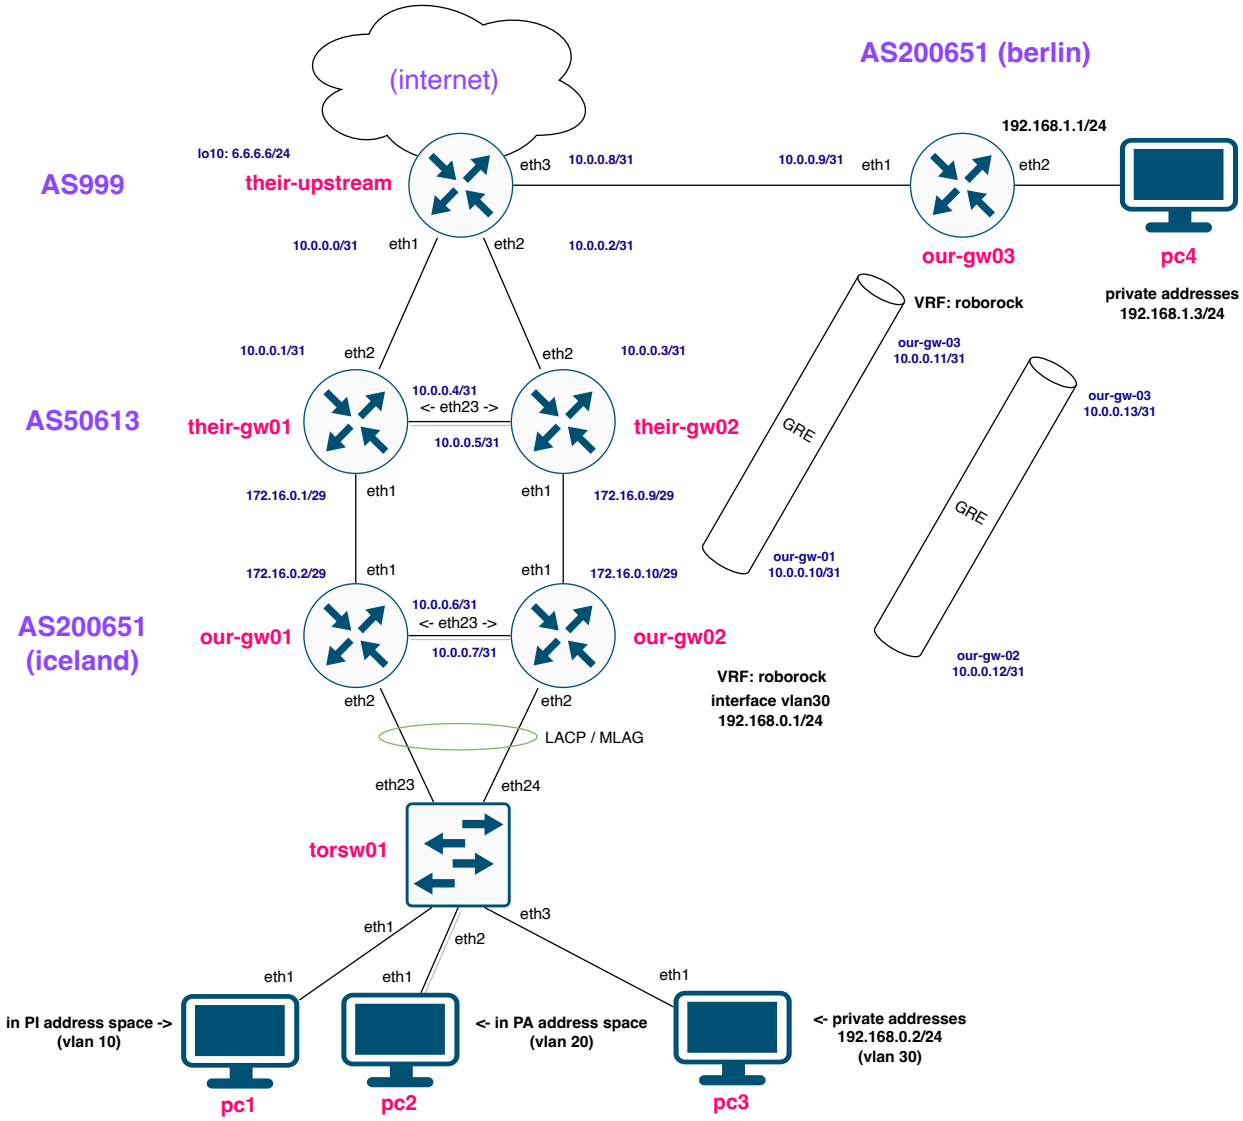

Goal

- Discuss different options for advertising PI space while also using PA space
- Implement them and run tests

Design goals

- our-gw-01 and our-gw-02 run eBGP with the provider
- our-gw-01 and our-gw-02 run iBGP with one another
- our-gw-01 and our-gw-02 provide a redundant default gateway for all customer networks
- torsw-01 should connect with LACP to the MLAG on our-gw-01 and our-gw-02
- we should be able to use both PI and PA address space

Discussion

- idea 1:
  - we announce both ranges from our AS
  - they use a route-map to delete our AS from that prefix
- idea 2:
  - they extend a VLAN to us containing the PA prefix
  - they would hold the gateway address (and hopefully provide redundancy)
  - how do we solve the two-routers problem?
- idea 3:
  - fancy routing magic
  - they route the PA prefix to an IP on our network
  - we advertise that IP to them over BGP
  - this would need a recursive lookup, maybe not possible? maybe? let's try
- idea 4:
  - two BGP sessions (one private, one public)
- idea 5:
  - some other tunneling protocol (GRE?)
- idea 6:
  - use neighbor local-as to pretend to be the private AS

Implementation 01

- configure `their-upstream` with:
  - interface loopback10, ip address 6.6.6.6/24
  - interface eth1, ip address 10.0.0.0/31 <- link to downstream
  - interface eth2, ip address 10.0.0.2/31 <- link to downstream
  - bgp AS 999, neighbor 10.0.0.1 remote-as 50613
  - bgp AS 999, neighbor 10.0.0.3 remote-as 50613
  - bgp AS 999, network 6.6.6.0/24

Implementation 02

- configure `their-gw01` with:
  - interface eth2, ip address 10.0.0.1/31 <- link to upstream
  - interface eth23, ip address 10.0.0.4/31 <- link to their-gw02
  - interface eth1, ip address 172.16.0.1/29 <- link to us
  - bgp AS 50613, neighbor 10.0.0.0 remote-as 999
  - bgp AS 50613, neighbor 10.0.0.5 remote-as 50613
  - bgp AS 50613, neighbor 172.16.0.2 remote-as 200651
- configure `their-gw02` with:
  - interface eth2, ip address 10.0.0.3/31 <- link to upstream
  - interface eth23, ip address 10.0.0.5/31 <- link to their-gw01
  - interface eth1, ip address 172.16.0.9/29 <- link to us
  - bgp AS 50613, neighbor 10.0.0.2 remote-as 999
  - bgp AS 50613, neighbor 10.0.0.4 remote-as 50613
  - bgp AS 50613, neighbor 172.16.0.10 remote-as 200651

Implementation 03

- configure `our-gw01` with:
  - vlan 10, vlan 20
  - vlan 4094 in trunk group mlagpeer
  - interface eth23, switchport mode trunk, switchport trunk group mlagpeer
  - interface vlan 4094, ip address 10.0.0.6/31
  - mlag, local-interface vlan 4094, peer-address 10.0.0.7, peer-link eth23, domain-id mlag1
  - interface eth1, ip address 172.16.0.2/29 <- link to upstream
  - interface eth2, channel-group 2 mode active
  - interface port-channel 2, mlag 2, switchport mode trunk, allowed vlan 10,20
  - interface vlan10, ip address 185.165.170.252/24
  - interface vlan20, ip address 82.221.100.252/23
  - bgp AS 200651, neighbor 172.16.0.1 remote-as 50613
  - bgp AS 200651, neighbor 10.0.0.7 remote-as 200651

Implementation 04

- configure `our-gw02` with:
  - vlan 10, vlan 20
  - vlan 4094 in trunk group mlagpeer
  - interface eth23, switchport mode trunk, switchport trunk group mlagpeer
  - interface vlan 4094, ip address 10.0.0.7/31
  - mlag, local-interface vlan 4094, peer-address 10.0.0.6, peer-link eth23, domain-id mlag1
  - interface eth1, ip address 172.16.0.10/29 <- link to upstream

  - interface eth2, channel-group 2 mode active
  - interface port-channel 2, mlag 2, switchport mode trunk, allowed vlan 10,20
  - interface vlan10, ip address 185.165.170.253/24
  - interface vlan20, ip address 82.221.100.253/23
  - bgp AS 200651, neighbor 172.16.0.9 remote-as 50613
  - bgp AS 200651, neighbor 10.0.0.6 remote-as 200651

Implementation 05

- configure `torsw01` with:
  - vlan 10, vlan 20
  - interface 23-24, channel-group 23 active
  - interface port-channel 23, switchport mode trunk, allowed vlan 10,20
  - interface eth1, switchport mode access, switchport access vlan 10
  - interface eth2, switchport mode access, switchport access vlan 20

Implementation 06 - VARP

- on our-gw01
  - ip virtual-router mac-address 0000.0000.0001
  - interface vlan 10, ip virtual-router address 185.165.170.254
  - interface vlan 20, ip virtual-router address 82.221.100.254
- on our-gw02
  - ip virtual-router mac-address 0000.0000.0001
  - interface vlan 10, ip virtual-router address 185.165.170.254
  - interface vlan 20, ip virtual-router address 82.221.100.254

Implementation 07 - PC IP addresses

- configure `pc1` with:
  - ip address add 185.165.170.1/24 dev eth1
  - ip route delete default
  - ip route add default via 185.165.170.254
- configure `pc2` with:
  - ip address add 82.221.100.1/24 dev eth1
  - ip route delete default
  - ip route add default via 82.221.100.254

Testing

- everything's up
  - ping 6.6.6.6 from pc1
  - ping 6.6.6.6 from pc2
- network breakage (loss of primary uplink)
  - shutdown our-gw01:eth1
  - ping 6.6.6.6 from pc1
  - ping 6.6.6.6 from pc2
  - no shutdown our-gw01:eth1
- network breakage (loss of TOR switch uplink)
  - shutdown our-gw01:eth2
  - ping 6.6.6.6 from pc1
  - ping 6.6.6.6 from pc2
  - no shutdown our-gw01:eth2
- network breakage (loss of secondary uplink)
  - shutdown our-gw02:eth1
  - ping 6.6.6.6 from pc1
  - ping 6.6.6.6 from pc2
  - no shutdown our-gw02:eth1

Add a GRE tunnel

- configure our-gw03's eth1 and eth2
- create a loopback interface with a flokinet IP, 185.165.171.1/24
- configure their-upstream's eth3
- establish BGP between their-upstream and our-gw03
- announce 185.165.171.0/24
- ensure that our prefixes are visible (`show bgp` for help if there are missing prefixes)
  - on our-gw-03
  - and on our-gw-01 and our-gw-02
- create a GRE tunnel from our-gw-01 to our-gw-03
- create a GRE tunnel from our-gw-02 to our-gw-03
- configure pc3 and pc4 (IP addresses and gateways)