S1E16 - The Border Gateway Protocol - Lab 02

Nicholas Morrison nick@nanocat.net

Connecting

Connect to the lab server:

```
local$ ssh-keygen -R netlab.nanocat.net
local$ ssh lab@netlab.nanocat.net
Password: (see discord)
```

Connect to your router:

```
lab@netlab$ list-devices
lab@netlab$ connect DEVICE
```

Introduction: About BGP ASNs and PI address space allocations

- Your RIR can allocate you your own:
  - an AS number (ASN)
  - IPv4 Provider-Independent (PI) address space
  - IPv6 Provider-Independent (PI) address space
- These allocations are GLOBAL (Internet-wide).
- There is a limited amount of address space, and a limited number of ASNs.
- Each RIR has its own policies for allocation. For example, https://www.ripe.net/manage-ips-and-asns/as-numbers/request-an-as-number
- If you have your own ASN and PI address space, you can ask your Internet Service Provider (ISP) to run a BGP session with you
- If you do not have your own ASN and PI address space, you can use a private ASN and borrow addresses from your ISP

Our new customers

- The customers in our new topology are divided into two groups:
  - those with their own ASNs and PI addresses
  - those without
- You will help customers with their own ASN and PI set up BGP peering sessions with you
- For customers without their own ASN, you will assign a private one, and carve out some of your own address space for them

Topology

Prefix and AS number allocations: https://docs.google.com/spreadsheets/d/1vxVgJI_LhY3NDGJ3icE-4zphmBxUurEvVFgdjafFBzc/edit?usp=sharing

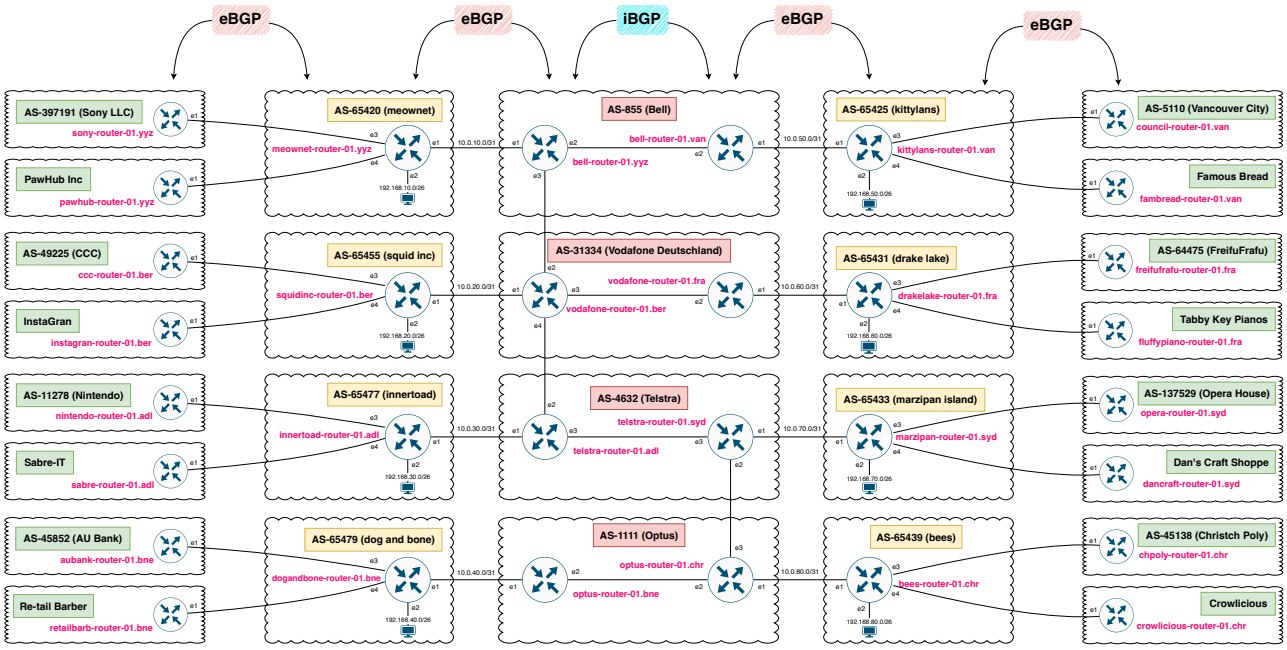

Open this image

Goal

- Find your customers in the spreadsheet
- Allocate ASNs and address space for customers without their own
- Configure BGP sessions for both of your customers
- Verify that it's all working

Allocate ASNs and address space

- Open the spreadsheet
- For the customers without their own ASN and PI address space:
  - Allocate an ASN from the 32-bit private ASN range
  - Allocate the next available /31 for the link between you and them
  - Allocate a /26 from your /24
    - Remember that you already allocated the first /26 for your PC!

Configure links

- Configure a /31 on their uplink to you
  - Customer gets the lower of the two addresses
- Configure the /31 on your downlink to them
  - You get the higher of the two addresses

Configure customer loopback interfaces

- Configure a new loopback interface on each customer router
  - loopback10
  - use the first IP address from their allocation
  - use the netmask from their allocation

Configure BGP sessions

- On your router, configure two new BGP sessions
  - one for each customer
  - use their /31 link address as the neighbour address
  - use the correct ASN for their remote-as
- On the customer's routers, configure BGP sessions
  - use the /31 link address as the neighbour address
  - use your own ASN for their remote-as
  - advertise their network with a network statement

Verify and Troubleshoot

- From your PC (which you will need to configure again):
  - traceroute to your two customers' loopback10 addresses
  - traceroute to other customers' loopback10 addresses

BONUS ROUND

- Level 1: Steal someone else's prefix and route it to yourself
- Level 2: Have your customer steal someone else's customer's prefix
- Level 3: Stop your customer from stealing someone else's customer's prefix

Appendix: configuration snippets

Configure an IP address

```
!
interface ethernet2
   no switchport
   ip address x.x.x.x/xx
!
```

Create a loopback interface

```
!
interface loopback10
   ip address x.x.x.x/xx
!
```

Configure a BGP session

```
!
ip routing
!
router bgp XXXX                        <- your ASN
   neighbor x.x.x.x remote-as YYYY     <- neighbor ASN
   network x.x.x.x/xx                  <- announce this network
!
```

For a network to be **announced** by BGP, an exactly matching route needs to exist in the routing table.

For a network to be **accepted** by BGP, the “BGP next-hop” address must be reachable.

Configure your PC

```
ip address add 192.168.XX.10/26 dev eth1
ip route delete default
ip route add default via 192.168.XX.1
```