# Peering automations using PDB

Antti Ristimäki, nog.fi meeting 2025.06 Tampere



A SUSTAINABLE FUTURE THROUGH DIGITALISATION

# Agenda

 An example how to automate (public) peering configs using PeeringDB as a source of truth





- Issues with manual peering configuration
  - Inconsistent configurations, naming conventions etc.
  - Obsolete peer descriptions (networks changing their names etc.)
  - Arbitrary and/or obsolete prefix-limits
  - Disconnected peers still lurking in the configs
- Automation is usually applied to bulk stuff  $\rightarrow$  peerings are just that
  - Design once, use repeatedly



### Elisa peering automation

- AS6667 public (IX) peerings migrated to PeeringDB based automated configs during 2024
- Ansible Jinja2 template, feeded with data from PeeringDB
- The Ansible playbook is not (yet) ran automatically, but instead triggered by an operator
  - Automatic GHA based config file background rendering, though
  - Some plans to deploy fully automated config push to the network



# Introducing an IX



an internal blob for IX identification

a pointer to the IX PeeringDB entry

list of routers connected to the IX

list (of dicts) of peer ASNs



#### Peer MD5 definitions

#### peers:

- { asn: xxxx, md5: \$9\$foobar }
- { asn: yyyy, ipv4\_md5: \$9\$foobar, ipv6\_md5: \$9\$barfoo }



# Non-selective peering with everybody

• A keyword 'all' can be given in the peer list, in which case all the peers that have defined their presence in the given IX in PeeringDB will be configured for the given IX

ixps:
<pre>- name: foobar_ix1</pre>
peeringdb_id: xxxxx
routers:
<pre>- magnificent_router1</pre>
peers:
– 'all'



# **Relevant PeeringDB objects**

- **net** object represents a network
- ix object represents an Internet Exchange
- ixlan object represents an IX peering LAN
  - it seems that ixlan\_id == ix\_id, but I'm not sure if this is always the case(?)
- netixlan object is a logical link between a network and the IX lan
  - one could say that a network is connected to an IX fabric via an netixlan object



### PeeringDB objects relationships

antti@A-WCP9PW0M2J ~ % curl -s -H "Authorization: Api-Key \$APIKEY" 'https://www.peeringdb.com/api/netixlan?asn=6667&fields=ixlan\_id,ix\_id,net\_id,name,asn,ipaddr4,ipaddr6' | jq



#### PeeringDB related Ansible tasks

- For each router to be configured, two PeeringDB API queries are initiated
- Only netixlan and net objects currently needed to render the config in our excercise
- As much information packed into single queries as possible ("ixlan\_id\_\_in" and "asn\_\_in") to avoid throttling

Get all **netixlan** objects from all IXes that the given router is connected to

For all the retrieved netixlan objects, get the respective **net** objects

- name: "Get netixlan information from PeeringDB"	
ansible.builtin.uri:	
<pre>url: "https://www.peeringdb.com/api/netixlan?ixlan_idin={{ ixp_list   join(',') }}"</pre>	
headers:	
Authorization: "Api-Key {{ peeringdb_api_key }}"	
timeout: 60	
ivars:	
<pre>ixp_list: "{{ ixps   default([])   selectattr('routers', 'search', inventory_hostname)   map(attribute='peeringdb_id')   unic</pre>	que }}"
register: peeringdb_netixlans	
when: ixp_list	
check_mode: False	
tags: always, ixp	
- name: Get IX peers information from Peeringdb	
ansible.builtin.uri:	
url: "https://www.peeringdb.com/apiźnet.json?asn_in={{ peer_asns   join(',') }}"	
headers:	
Authorization: "Api-Key {{ peeringdb_api_key }}"	
timeout: 60	
tvars:	
<pre>peer_asns: "{{ peeringdb_netixlans.json.data   map(attribute='asn')   unique }}"</pre>	
'register: peeringdb_1x_peers	
when: ixps   default([])   selectattr('routers', 'search', inventory_hostname)   map(attribute='peeringdb_id')   list	
check_mode: False	
tags: always, ixp	

# **Rendered configuration**

```
protocols {
 bgp {
    replace:
    /* configured by Ansible, please do not edit manually */
    group ficix1_ipv4 {
      description "FICIX 1 (Espoo)";
      type external;
      local-address 193.110.226.19;
      remove-private;
      import [our import policy chain];
      export [ our export policy chain];
      neighbor 193.110.226.14 {
        description "PEER - FICIX 1 (Espoo); CSC/Funet";
        family inet {
          unicast {
            prefix-limit {
              maximum 100;
              teardown {
                 90;
                 idle-timeout 60;
        authentication-key "$9$XXXXXXXXXXXXXX;;
        peer-as 1741;
      neighbor 193.110.226.xx {
```

eliso

# About PeeringDB prefix-limits

- It is expected that the prefix count in PeeringDB represents the suggested max-prefixes setting
  - IMHO the order of magnitude is more relevant that the exact number of advertised prefixes
- In our template a PeeringDB-found value <10 is always rounded to a value of 10 and a hard-coded default value of 1000 is used if no value in PeeringDB

{% if peer\_dict.info\_prefixes4 is defined and peer\_dict.info\_prefixes4 > 0 %}
maximum {{ [ 10, peer\_dict.info\_prefixes4 ] | max }};
{% else %}
maximum 1000;
{% endif %}

Recommended maximum number of IPv4 routes/prefixes to be configured on peering sessions for this ASN. Leave blank for not disclosed.

IPv4 Prefixes ??

IPv4 Prefixes 🕐		7
IPv6 Prefixes 🕜		10
v4 Prefixes 🕜		1
IPv6 Prefixes 🕐	1	1
IPv4 Prefixes ?	564	
IPv6 Prefixes 2	64	



# PeeringDB API throttling

- Even authenticated PeeringDB API queries are rate-limited (40 queries/min)
- Possible solutions
  - keep your own local mirrored cache
  - pack more stuff into a single query
  - partition your workflow into smaller batches and pause in between
- In our case, the Ansible playbook is usually ran against a single router at a time, so ratelimiting is not a big issue
- We also try to pack as much information into a single query as possible



#### **Route servers**

- Excluded from automation
- After all, peering with route servers could be considered an alternative to bilateral peering
  - Exclusively peering with route servers  $\rightarrow$  no need to automate bilateral peerings
  - Bilateral peerings automated  $\rightarrow$  less incentives to peer with route servers



#### Future work

- (At least semi-)Automate messaging with peers
- Better management of MD5 (and/or TCP-AO) secrets
- Somehow better manage configured yet non-established sessions
- Include IRR based prefix filtering



