Communication Operator(KO)

Keeping 1k+ IP pools filled with just in time provisioning

Kristoffer Larsen

Data Network Design Expert

kri@telia.net

Telia Company

Current Design of KO within Telia Company

- Started around 2004 with different models of 100Mb Ethernet access switches. Some of these models are still in use today.
- Built around a service model “copied” from ADSL where QoS was based on multiple Sessions on PVCs or in KO’s case VLANs.
- Most of the service intelligence is placed in the access node.
- Each service provider will have one or more VLANs in the access network.

Change Motivation

- Service design became increasingly complicated and/or differentiated with 40+ different access switch models.
- Lack of IPv6 support.
- More and more service providers meant a large deployment task to provision VLANs - and to some extent older switches could not support the amount of VLANs/service providers needed.
- Much manual work involved in keeping service provider IPv4 pools full on 300 service edges for 60 service providers.

New Design

- Extensive use of 802.1q QinQ.
- Each customer has their own VLAN towards the BNG.
- Simpler access nodes with limited L3 awareness, except for IGMP that is still needed for multicast.
- Less and larger redundant BNG pairs, from 300 to 60-70 service edges.
- Support for native IPv6.
- Single UNI allowing for free-seating.

Our Challenges with IPv4 Address Exhaustion

- There is a demand for very effective address usage.
- Customers can move between service providers 24/7.
- 2 different solutions based on who manages the DHCP
    - Service provider managed.
    - Communication operator managed.

Service Provider managed DHCP

- BNG learn /32 host routes based on answers from DHCP servers.
- Proxy ARP forces all traffic in the BNG direction.
- A static subnet is agreed for gi-addr, can be private. Each BNG uses a /32 from this subnet.
- In order to scale subscribers, FIB loading is limited to NNI routers and their BNG pair.
- We provide no aggregation of routes. If desired it needs to be done on the service provider side.
- All upstream traffic is routed directly to the NNI, including the gateway provided from the server.

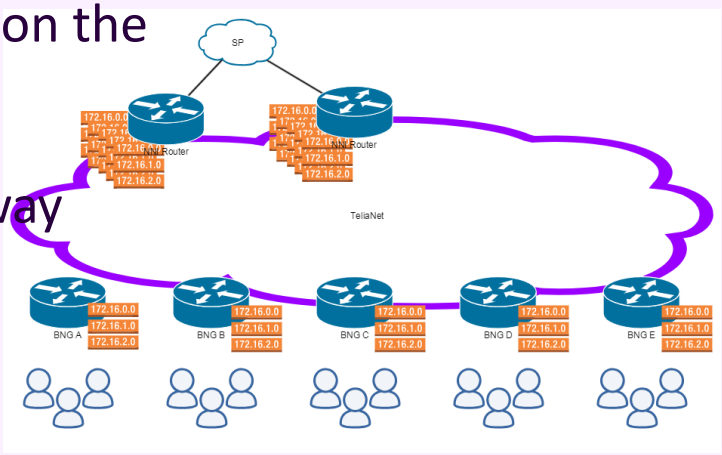

Communication Operator managed DHCP

- Each service provider has a large public and/or CGNAT pool for the entire country.
- This large and centralized pool will be subnetted into suitable BNG-chunks.
- Initial state of the BNG is that no IPv4 addresses are assigned.
- Based on demand, pools will be added to the BNG just in time.

Assignment Process

- When a pool is missing or is saturated the BNG sends a NETCONF notification
- The orchestration system streams to these notifications.
- They are nicely formatted with clear indication of message type, and key fields separated.
- From this notification, we locate the exact pool and its assignment size.
- Multiple DHCP requests are handled using throttling while we are adding addresses.
- If all checks pass, we create an assignment request and apply the returned address to the BNG.
- Process normally takes around 4 seconds.

```
<notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
  <eventTime>2022-09-06T13:11:49.657+02:00</eventTime>
  <sros-log-generic-event xmlns="urn:nokia.com:sros:ns:yang:sr:notifications">
    <sequence-number>1104</sequence-number>
    <severity>warning</severity>
    <application>dhcps</application>
    <event-id>2003</event-id>
    <event-name>tmnxDhcpSvrPoolUnknown</event-name>
    <router-name>vprn65</router-name>
    <subject>Unknown pool</subject>
    <message>DHCP server TLA_V4 detects an unknown pool (sbg-ref-sec21:TLA-1).  Pool extracted fro
    <event-params>
      <tmnxDhcpSvrNotifVRtrID>16</tmnxDhcpSvrNotifVRtrID>
      <tmnxDhcpSvrNotifServerName>TLA_V4</tmnxDhcpSvrNotifServerName>
      <tmnxDhcpSvrNotifUnknownPoolName>sbg-ref-sec21:TLA-1</tmnxDhcpSvrNotifUnknownPoolName>
      <tmnxDhcpSvrNotifDescription> Pool extracted from dhcp-message is unknown in the server.</tm
      <tmnxDhcpSvrNotifMsgHwAddress>00:16:01:00:00:07</tmnxDhcpSvrNotifMsgHwAddress>
      <tmnxDhcpSvrNotifClientDUID/>
    </event-params>
  </sros-log-generic-event>
</notification>
```

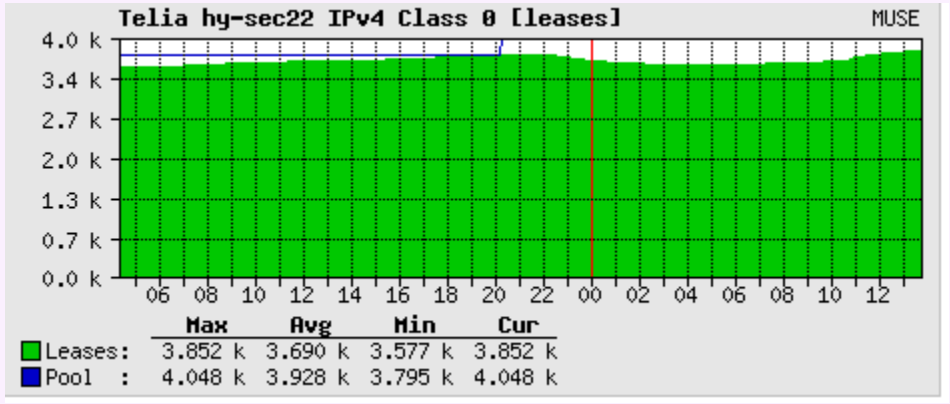

Returning Addresses

- Due to network optimization or customers changing service provider, there is a constant demand to also be able shrink pools.
- Currently a semi automated process is used for this.
- Process
    - Initially tell the DHCP server on the BNG to stop renewing addresses on the specific subnet.
    - Start monitoring the stable lease state of the subnet.
    - Once the subnet is completely empty, we remove it from DHCP and routing.
- We always take the least used subnet, but it may result in a small impact on customers when changing their IPv4. E.g. TCP or VPN connections are lost and then have to be re-started.
- We are planning to automate this, but the framework is still under discussion, things like weekly peak usage and available addresses in the overall pool. We aim to only do it when needed to limit customer impact.

Coming Steps

- Migration into this new design has started, some of you might know it as GEKO.
- We are at around 20% now hoping to be much further by the end of next year.
- Which gives us more than another million ports to migrate.

Questions

?