

Case Study: Renumbering an AS

Pim van Pelt <pim@ipng.ch> • 2021-11-23 • RIPE #83



Introduction



Pim van Pelt

Pim van Pelt (PBVP1-RIPE)

- Member of the RIPE community since 1999 (RIPE #34)
 - Has used pim@ipng.nl for 22 years
 - And also pim@ipng.ch for 15 years
 - Incorporated ipng.ch in Switzerland in 2021





Introduction



IPv6 Deployment & Tunnel Broker




[Enter the SixXS website](#)

SixXS - a global IPv6 Tunnelbroker

- Member of the RIPE community since 1999 (RIPE #34)
 - Founded SixXS w/ Jeroen Massar [[ref](#)]
 - Operated SixXS from ca. 1999 - 2017 [[sunset](#)]
 - Used **AS8298** for Ghost Route Hunter (2003) [[RIPE #44](#)]
-



Introduction





IPng Networks GmbH

- Developer of Software Routers - VPP and DPDK [[ref](#)]
 - Tiny operator from Zurich, Switzerland [[ref](#)]
 - Acquired AS8298 from SixXS
-



Introduction



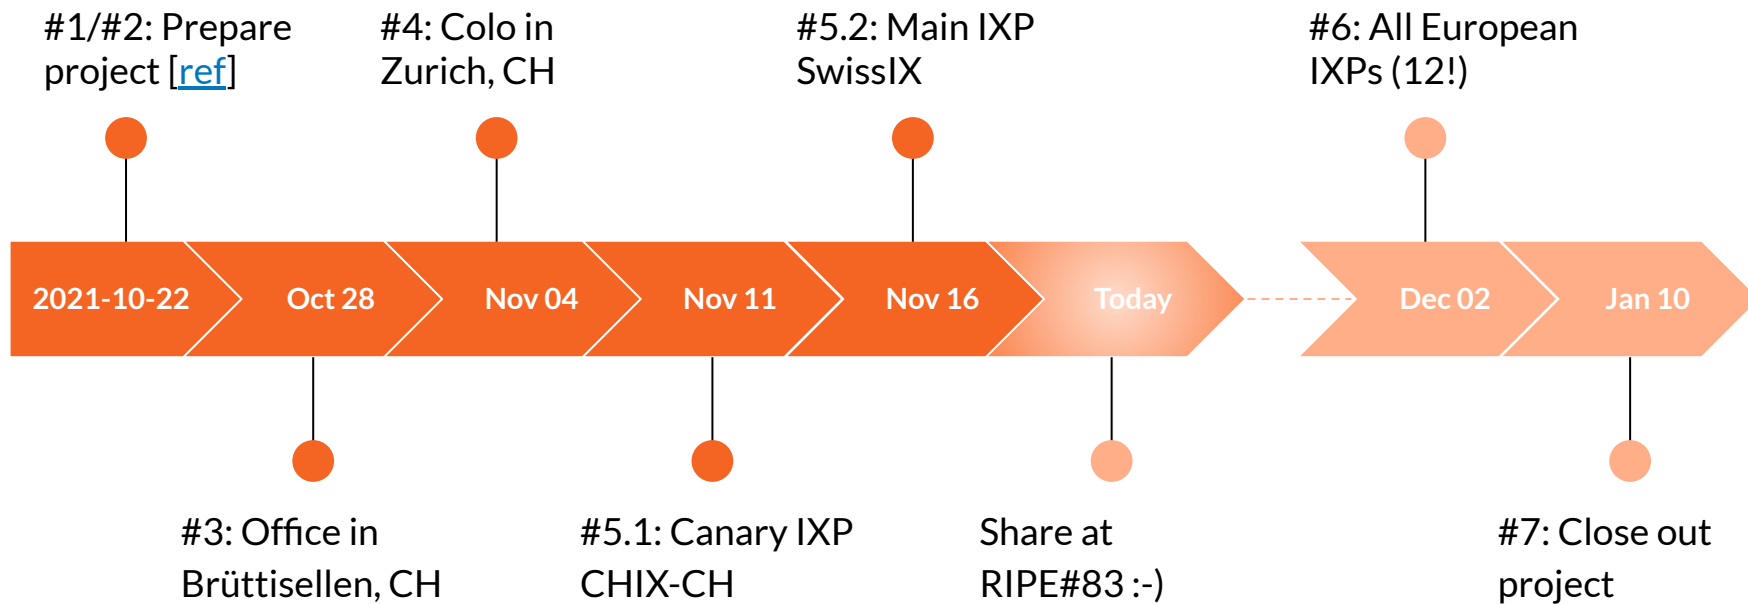
IPng Networks GmbH

- Developer of Software Routers - VPP and DPDK [[ref](#)]
 - Tiny operator from Zurich, Switzerland [[ref](#)]
 - Acquired AS8298 from SixXS
 - Twelve VPP/Bird2 routers [[ref](#)] (UN/LOCODE names)
 - European ring: *peering on the FLAP** [[ref](#)] ~1850 adjacencies
-



Project Timeline

2021-10-22 - 2022-01-10






Preparation



Step #0 - Acquire the AS

1. Ask original + current holder
 2. Transfer agreement
 3. Sponsoring LIR handoff
 4. RIPE NCC paperwork
 5. Copy route, route6 and RPKI ROA from AS50869 to AS8298
-




Preparation

Step #1 + #2 - Split AS50869

1. Distribute connected and static/blackhole into OSPF
2. Restrict routers that originate (from **all** to only **chbt10**, **chbt11**)

... no changes visible in the DFZ yet




Step #3 - Convert our Office

Announce 3x IPv4 and 3x IPv6 from AS8298 to AS50869

Topology changes:

- **chbt10** (8298) \Leftarrow **chrma0.ipng.ch** (50869)
- **chbt11** (8298) \Leftarrow **chgtg0.ipng.ch** (50869)
- Create **chbt10**, **chbt11** in iBGP (8298)
- LIR ch.ipng is now ^50869_8298\$
- 8298 also connects to 58299, 58280 in Brüttisellen, CH
 - Thanks to 58299 for local IP transit!



Step #4 - Convert our Colo

Put first customer AS57777 behind AS8298

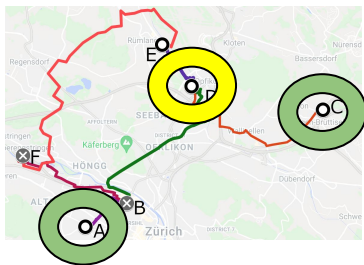
Topology changes:

- **ddl1n0** (8298) \Leftarrow **chgtg0.ipng.ch** (50869)
- **ddl1n1** (8298) \Leftarrow **chrma0.ipng.ch** (50869)
- Add **ddl1n0, ddl1n1** to iBGP (8298)
- LIR **ch.massar** (colo customer) is now **^50869_8298_57777\$**

Notable:

- Colo is reachable from **all** routers due to OSPF *connected*
- Even though BGP is AS8298, it's reachable via AS50869
- So this is a stable situation





Step #5.1 - Connect our office to our colo

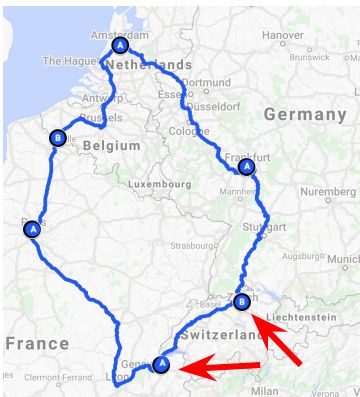
Connect the two islands **chbt1** and **ddl1n** via **chgtg0.ipng.ch**

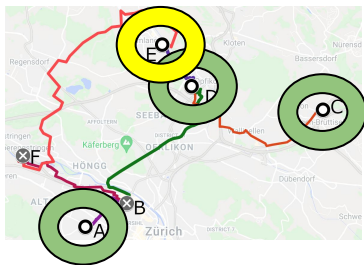
Topology changes:

- **chgtg0** (8298) \Leftarrow **chrma0.ipng.ch** (50869)
- **chgtg0** (8298) \Leftarrow **chplo0.ipng.ch** (50869)
- **chbt10, chbt11** - **chgtg0** - **ddl1n0, ddl1n1** iBGP (8298)
- CHIX-CH Peers now see AS8298

Notable:

- Five routers in 8298 are now solidly connected to:
 - 50869 (2x), 58299 (2x) and 25091 (1x) transit
- Google 15169 and Init7 13030 most important peers
- There's ~32 broken BGP sessions at this point





Step #5.2 - Finish metro Zurich


Complete the Zurich Metro ring - all routers in AS8298


Topology changes:

- **chrma0** (8298) \leftarrow **defra0.ipng.ch** (50869)
- **chgtg0** (8298) \leftarrow **chp1o0.ipng.ch** (50869)
- Swiss peers now all see AS8298
 - CHIX-CH, Community IX, FreeIX and SwissIX

Notable:

- Six routers in 8298 are now redundantly connected to:
 - 50869 (2x), 58299 (2x), 25091 (2x) and 34549 (1x) transit
- There's ~**188 broken BGP sessions** at this point





Takeaways - automation is key (1/2)

```
pim@spongebob:~$ vim blobs/chplo0.ipng.ch/etc/bird/manual.conf
pim@spongebob:~$ vim blobs/defra0.ipng.ch/etc/bird/manual.conf
pim@spongebob:~$ vim config/defra0.ipng.ch.yaml
pim@spongebob:~$ vim config/chplo0.ipng.ch.yaml
pim@spongebob:~$ ROUTERS="chplo0.ipng.ch defra0.ipng.ch" ./build.sh


pim@spongebob:~$ git diff build/
pim@spongebob:~$ kees-push.sh chplo0.ipng.ch
pim@spongebob:~$ kees-push.sh defra0.ipng.ch

pim@spongebob:~$ git submit -m "AS8298 renumber step #5 (part 2) - prepare chplo0 and
defra0 (still in AS50869) to accept AS8298 as a downstream"
```

Inspired by, but a rewrite of, Coloclue's Kees [[github](#)]

- YAML configuration files
- Jinja2 templates for Bird, Unbound, Firewall, etc
- Ability to add manual configuration for esoteric scenarios






Takeaways - automation is key (2/2)

```
pim@chgtg0:~$ birdc show proto | awk -F_ '/chix.*BGP.*Bad peer AS/ { print $2 }' | sort -u | xargs
204438 213126 28859 48550 48646 51778

pim@spongebob:~$ send-email-by-asn --asn 204438 213126 28859 48550 48646 51778 --ixp 2365 \
--email emails/chix-as8298-1.j2
[INFO] main: AS numbers: [204438, 213126, 28859, 48550, 48646, 51778]
[INFO] main: PeeringDB IXP: id=2365 name='CH Swiss Internet Exchange'
[INFO] main: PeeringDB NetIDs: [1634, 9285, 12613, 22422, 17526, 23483]
[INFO] main: Sending e-mail: asn=28859 from=noc@ipng.ch to=noc@zugernet.ch, andreas.fries@acdalis.ch
[INFO] main: Sending e-mail: asn=48550 from=noc@ipng.ch to=noc@snapserv.net
[INFO] main: Sending e-mail: asn=48646 from=noc@ipng.ch to=noc@as48646.de
[INFO] main: Sending e-mail: asn=51778 from=noc@ipng.ch to=raphael.muther@pronetic.ch
[INFO] main: Sending e-mail: asn=204438 from=noc@ipng.ch to=noc@rubengarbade.ch
[INFO] main: Sending e-mail: asn=213126 from=noc@ipng.ch to=noc@afries.ch
```





Last Step #6 - Complete the ring!

Topology changes scheduled for 2021-12-02:

- AS8298 - **defra0, nlams0, frggh0, frpar0, chplo0**
- Establish iBGP on **rr0.{chbt10, frggh0, chplo0}.ipng.ch**
- 12 additional IXPs and ~1600 adjacencies to convert:

defra0 12552, 16839, 20940, 211620, 24429, 29169, 3856, 39572, 42, 43252, 45102, 46489, 47228, 56890, 6695, 6939, 714, 8966

nlams0 13335, 203729, 20495, 206238, 209288, 212635, 213035, 32934, 34307, 41441, 44103, 44854, 45009, 49917, 56393, 57866, 6939, 714, 8315

frggh0 16276, 16347, 200780, 2200, 34019, 35625, 47214

frpar0 20940, 24429, 29169, 32934, 37100, 3856, 42, 42064, 45102, 46489, 47406, 51706, 714, 8966


chplo0 12307, 35202, 39540, 44596, 48550, 51530, 51786, 57118

*Your AS here?
Please help me
renumber!*

Questions, Discussion

chbt102.paposting.net

Time range: to



1. If you peer with IPng Networks, thanks! Please help renumber :)
2. If you don't: please peer with AS8298 <peering@ipng.ch>
3. Success means:
 - a. No outages (so far, so good)
 - b. <100 adjacencies dropped

Also: thanks for listening!