

OpenBGPD



Claudio Jeker

OpenBGPD for IXP route-servers

OpenBGPD at IXP

Running OpenBGPD as a route-server with full filters is too slow

Old RIB design shows troubles with scaling

Config reload block updates and withdraw until finished

At a small IXP like YYCIX config requires 370'000 filter rules and reloading the config takes 1h to finish

As a result most IXP route-server run a bird monoculture

It is important to make OpenBGPD a viable alternative again

A major effort is required

Solving the issue

Have a full time developer working on OpenBGPD funded by



Community Projects Fund

and



OpenBGPD in 6.4

- RFC8212 compliance (default deny policy)
Be careful when updating from 6.3 to 6.4.
- RPKI ROA support (static table, no RTR support)
- Sets for prefixes, ASnum and origins (prefix + source-as)
Replace large amount of filters with a single fast lookup
- Background soft-reconfiguration on config reload
On reload new withdraws and updates are now processed
- 154 commits since 6.3 (close to 8% of all commits)

Results

YYCIX (Calgary Internet Exchange) using arouteserver to generate config

6.3: generated config consists of 370'000 filter rules

6.4: with as-set, prefix-set and origin-set ruleset is now below 6000 rules

Switching to as-set & prefix-set for IRR based filtering reduced initial config to 95'000 rules. Using origin-set for the RPKI and ARIN DB filtering did the rest.

Background soft-reconfiguration no longer blocks update processing during a config reload.

OpenBGPD in 6.5 (May 2019)

Better community filtering

- Matching more than one community per filter rule

- Faster set and delete of communities

More filter refinements and tuning

Further RIB refactoring, make multiple RIB support faster

Resurrect portable version and keep it alive (like OpenSSH and LibreSSL)

- No FIB support for now

Future Projects

Multi-thread support in the RDE for filtering and maybe some other tasks

In the portable version add FIB support for other BSDs and Linux

New features (ADD_PATH, BGP multipath, BMP, RTR)

Thanks



Community Projects Fund

and



Asteroid



also a big thank you goes to Job Snijders for making it possible

Questions?