

BIRD update RS memory optimization

Ondrej Filip / of@nix.cz
NIX.CZ

Sep 26, 2012, RIPE-65, EIX,
Amsterdam


Since last update

- ➔ Last update in Vienna (1.3.4) – extended community filters
- ➔ More deployments
- ➔ New key BGP related features
 - ⑤ Protocol templates
 - ⑤ ROA basic support
 - ⑤ Import/export route limits improved
 - ⑤ Secondary route export



BIRD deployments

- ➔ CDNs, IXPs – check Euro-IX report



ROA

- ➔ Support of ROA tables
- ➔ Can be dynamically filled (from RPKI or IRR db) using CLI
- ➔ ROA can be matched in filters -
roa_check()
 - ⑤ ROA_UNKNOWN
 - ⑤ ROA_VALID
 - ⑤ ROA_INVALID

```
roa table myroa {  
    roa 217.31.192.0/20 max 20 as 25192;  
}
```

Route limits

- ➔ General protocol option (not just for BGP)
- ➔ Number of routes received from or **sent to** a peer
- ➔ Four types of reaction
 - ⑤ WARN
 - ⑤ BLOCK
 - ⑤ RESTART
 - ⑤ DISABLE

Protocol templates

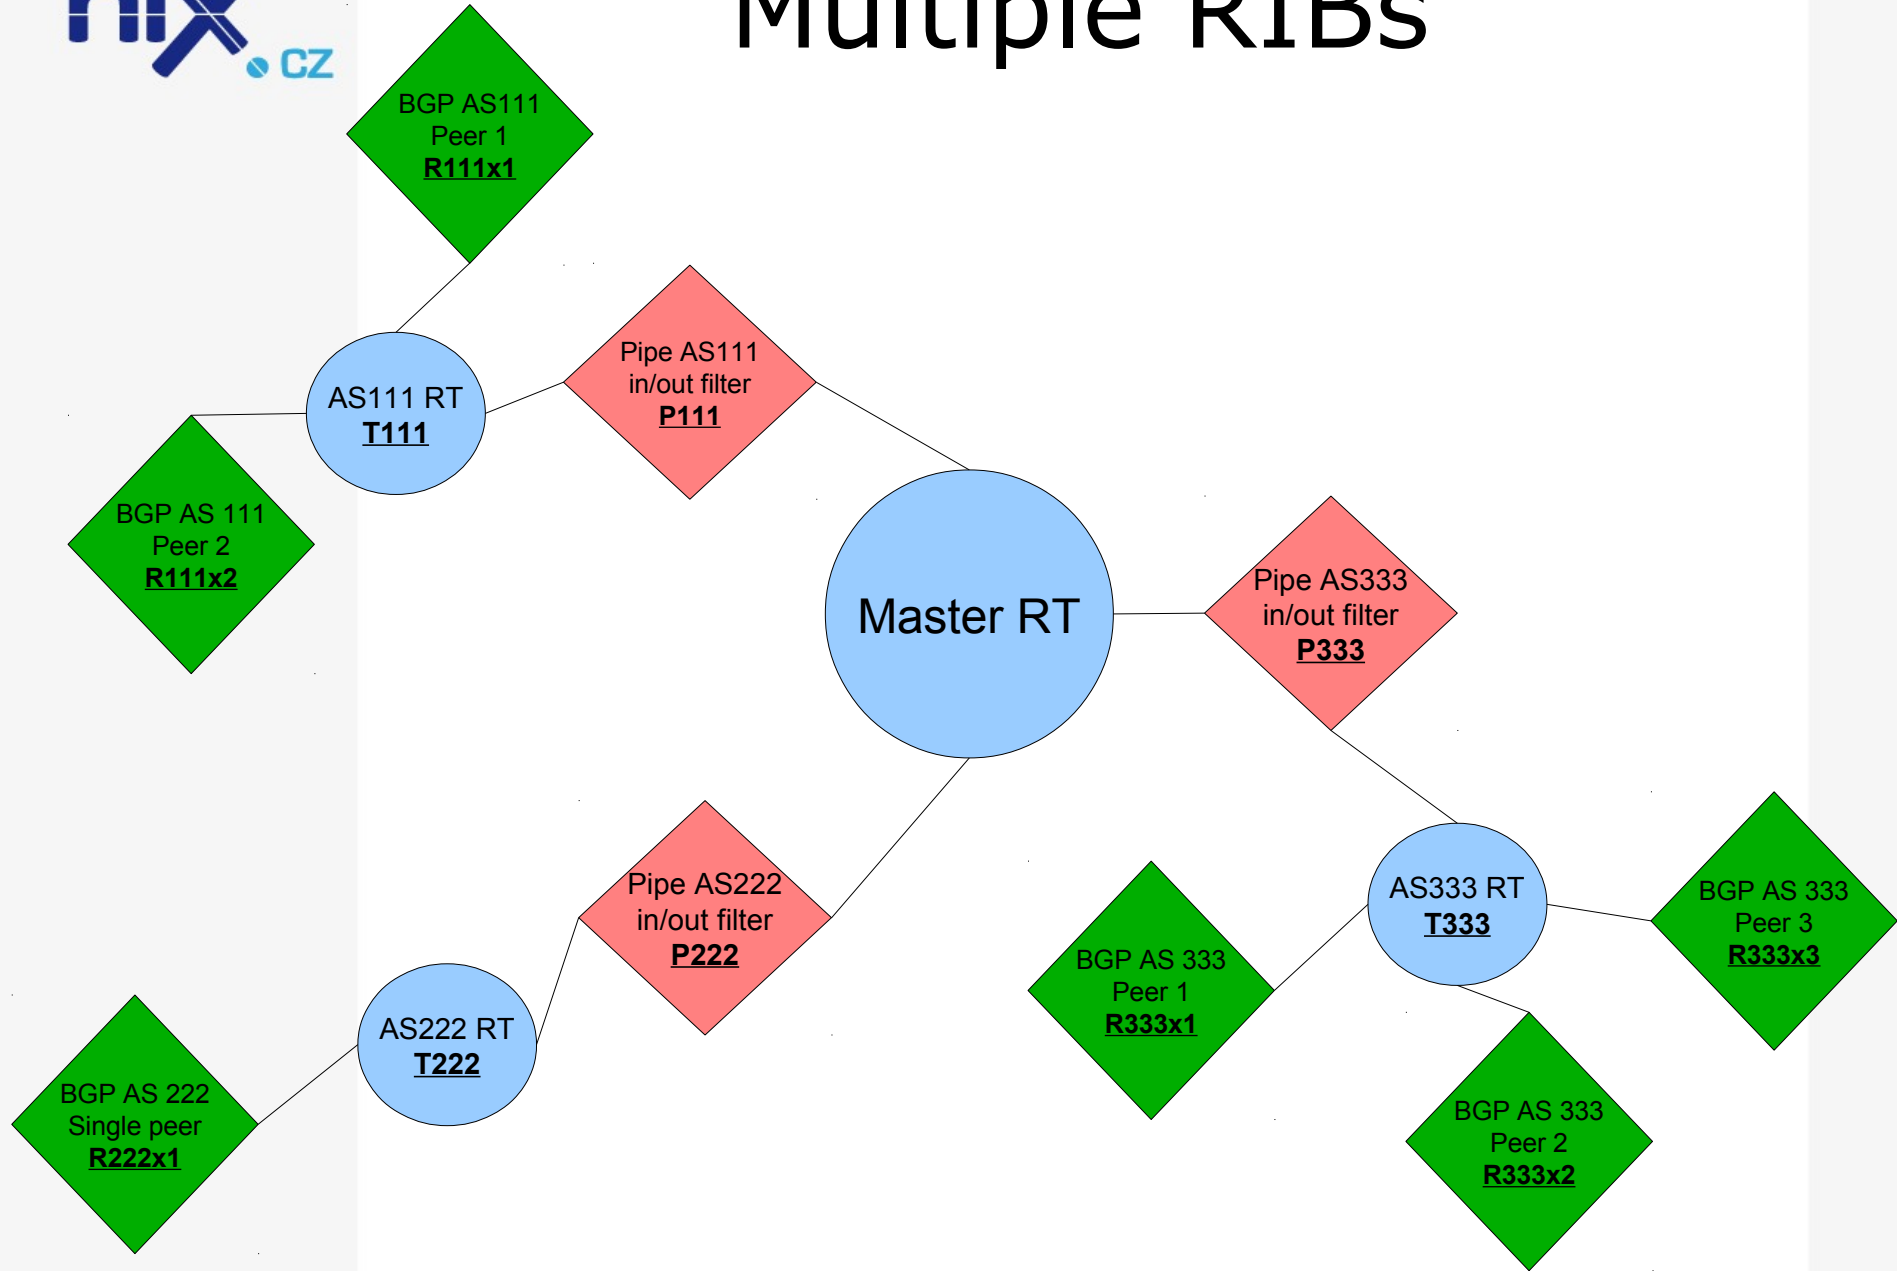
```
template bgp NIXPEERS {  
    local as 112;  
    export filter bgp_out;  
    start delay time 120;  
    debug { states, events };  
    route limit 50000;  
}
```

```
protocol bgp NIXRS1 from NIXPEERS {  
    neighbor 91.210.16.1 as 47200;  
    route limit 60000;  
}
```

Side note: Route filtering problem

- ➔ Normally filters match best route only
- ➔ If best route denied no other exported
- ➔ Problem with customers behind two different peering partners if one partner (with better path) is filtered out
- ➔ One solution multiple RIBs with PIPE protocol – but consumes a lot of memory (square complexity to # of routes)

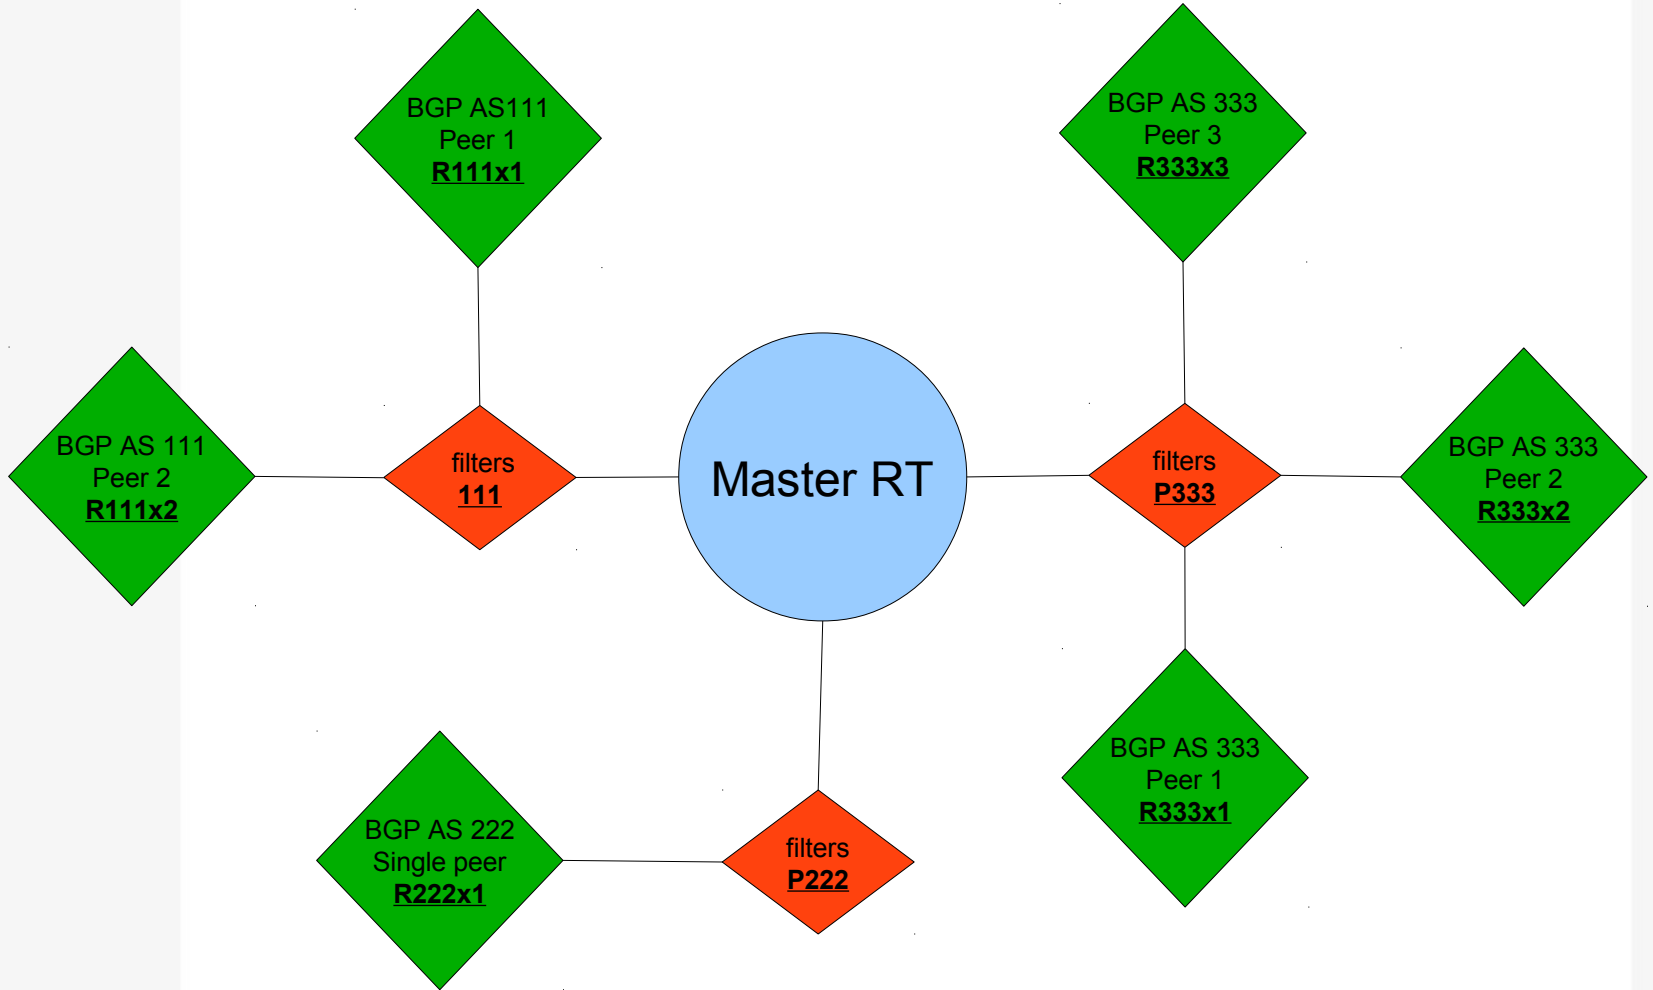
Multiple RIBs



Single RIB

- ➔ Allow second best route to be exported
- ➔ BGP protocol option 'secondary;'
- ➔ Precondition: routing table sorted (a little bit slower)
 - ⑤ 'table master sorted;'
- ➔ No PIPE protocols
- ➔ Con - Can't check imported routes before filtering

Single RIB





```
template bgp NIXPEERS {  
    local as myas;  
    rs client;  
    secondary;  
    route limit 25000;  
    start delay time 60;  
    interpret communities off;  
}
```

```
table master sorted;
```

```
protocol bgp R6881x1 from NIXPEERS {  
    description "NIX.CZ - 91.210.16.243 - (1)";  
    neighbor 91.210.16.243 as 6881;  
    import filter bgp_in_AS6881;  
    export where bgp_out(6881);  
}
```

NIX.CZ Route Servers

- ➔ Two route servers
- ➔ BIRD 1.3.8 with single RIB configuration on Linux
- ➔ BIRD 1.3.5 with multiple RIBs configuration on FreeBSD
- ➔ ~30.000 IPv4 routes
- ➔ ~7.000 IPv6 routes
- ➔ Memory consumption 1/7 IPv4 and 1/4 IPv6

Work in progress

- ➔ BGP Add-path (beta)
- ➔ Looking glass (beta)
- ➔ IPv4 and IPv6 integration
- ➔ IS-IS
- ➔ ... your feedback is welcome

Conclusion

- ➔ Secondary route export can save a lot of memory/CPU with
- ➔ More important for larger deployments
- ➔ New BGP features in BIRD
- ➔ And look forward to more :-)

Questions ?

<http://www.nix.cz>

<http://bird.network.cz>

<mailto:of@nix.cz>