



# BIRD Internet Routing Daemon

Update, Route Servers



Ondrej Filip • [ondrej.filip@nic.cz](mailto:ondrej.filip@nic.cz)  
06 Feb 2013 • NANOG 57



# CZ.NIC, CZ.NIC Labs



- Not just domain registry for .cz
- About 1M domains, 40% DNSSEC signed
- R&D department – CZ.NIC Labs
- Project for the good of the Internet
- DNSSEC Plugin, DANE patrol, Knot DNS, ...
- Check <http://labs.nic.cz>



# BIRD



- Internet Routing Daemon
- BGP, OSPF, RIP
- IPv4, IPv6, dual compilation
- Fast, efficient
- Powerful configuration and **filtering** language – variables, functions; binary trees - quick
- Multiple routing tables, PIPE protocol
- Introduced at NANOG 48



# Since NANOG 48 – major additions

- BGP
  - Improved community list matching
  - Extended communities + filtering
  - ROA basic support
  - Secondary route export
  - TTL security



# Since NANOG 48 – major additions

- IPv6 Router Advertisement
- “include” configuration option
- Re-configuration timeout
- OSPF NSSA
- Multipath support – OSPF and static
- Protocol templates
- Import/export route limits improved
- More deployments



# Deployed at ... (and more)



# Re-configuration timeout

```
bird> configure timeout 20
```

```
Reading configuration from /etc/bird.conf
```

```
Undo scheduled in 20 s
```

```
Reconfiguration in progress
```

```
bird> show protocols rip1
```

name	proto	table	state	since	info
rip1	RIP	master	down	18:46	

```
bird> >>> Config timeout expired, starting undo
```

```
show protocols rip1
```

name	proto	table	state	since	info
rip1	RIP	master	up	18:46	

```
.
```



# Re-configuration timeout

```
bird> configure timeout 20
```

```
Reading configuration from /etc/bird.conf
```

```
Undo scheduled in 20 s
```

```
Reconfiguration in progress
```

```
bird> show protocols rip1
```

name	proto	table	state	since	info
rip1	RIP	master	down	18:46	

```
bird> configure confirm
```

```
Reconfiguration confirmed
```

```
bird> show protocols rip1
```

name	proto	table	state	since	info
rip1	RIP	master	down	18:46	





# ROA support

- Support of ROA tables
- Can be dynamically filled (from RPKI or IRR db) using CLI or config file
- ROA can be matched in filters - roa\_check()
  - ROA\_UNKNOWN
  - ROA\_VALID
  - 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/protocol
- Four types of reaction
  - WARN
  - BLOCK
  - RESTART
  - DISABLE



# Protocol templates

```
template bgp NIXPEERS {  
    local as 112;  
    export filter bgp_out;  
    start delay time 120;  
    mrtdump all;  
    import limit 50000 action warn;  
}  
  
protocol bgp NIXRS1 from NIXPEERS {  
    neighbor 91.210.16.1 as 47200;  
    import limit 60000 action block;  
}
```



# Extended community lists

```
if peeras > 65535 then
{
  if (ro,0,peeras) ~ bgp_ext_community then return false;
  if (ro,myas,peeras) ~ bgp_ext_community then return true;
  if ((ro,0,myas) ~ bgp_ext_community) then return false;
} else {
  if ((0,peeras) ~ bgp_community) || ((ro,0,peeras) ~
      bgp_ext_community) then return false;
  if ((myas,peeras) ~ bgp_community) || ((ro,myas,peeras) ~
      bgp_ext_community) then return true;
  if ((0, myas) ~ bgp_community) || ((ro,0,myas) ~
      bgp_ext_community) then return false;
}
return true;
```

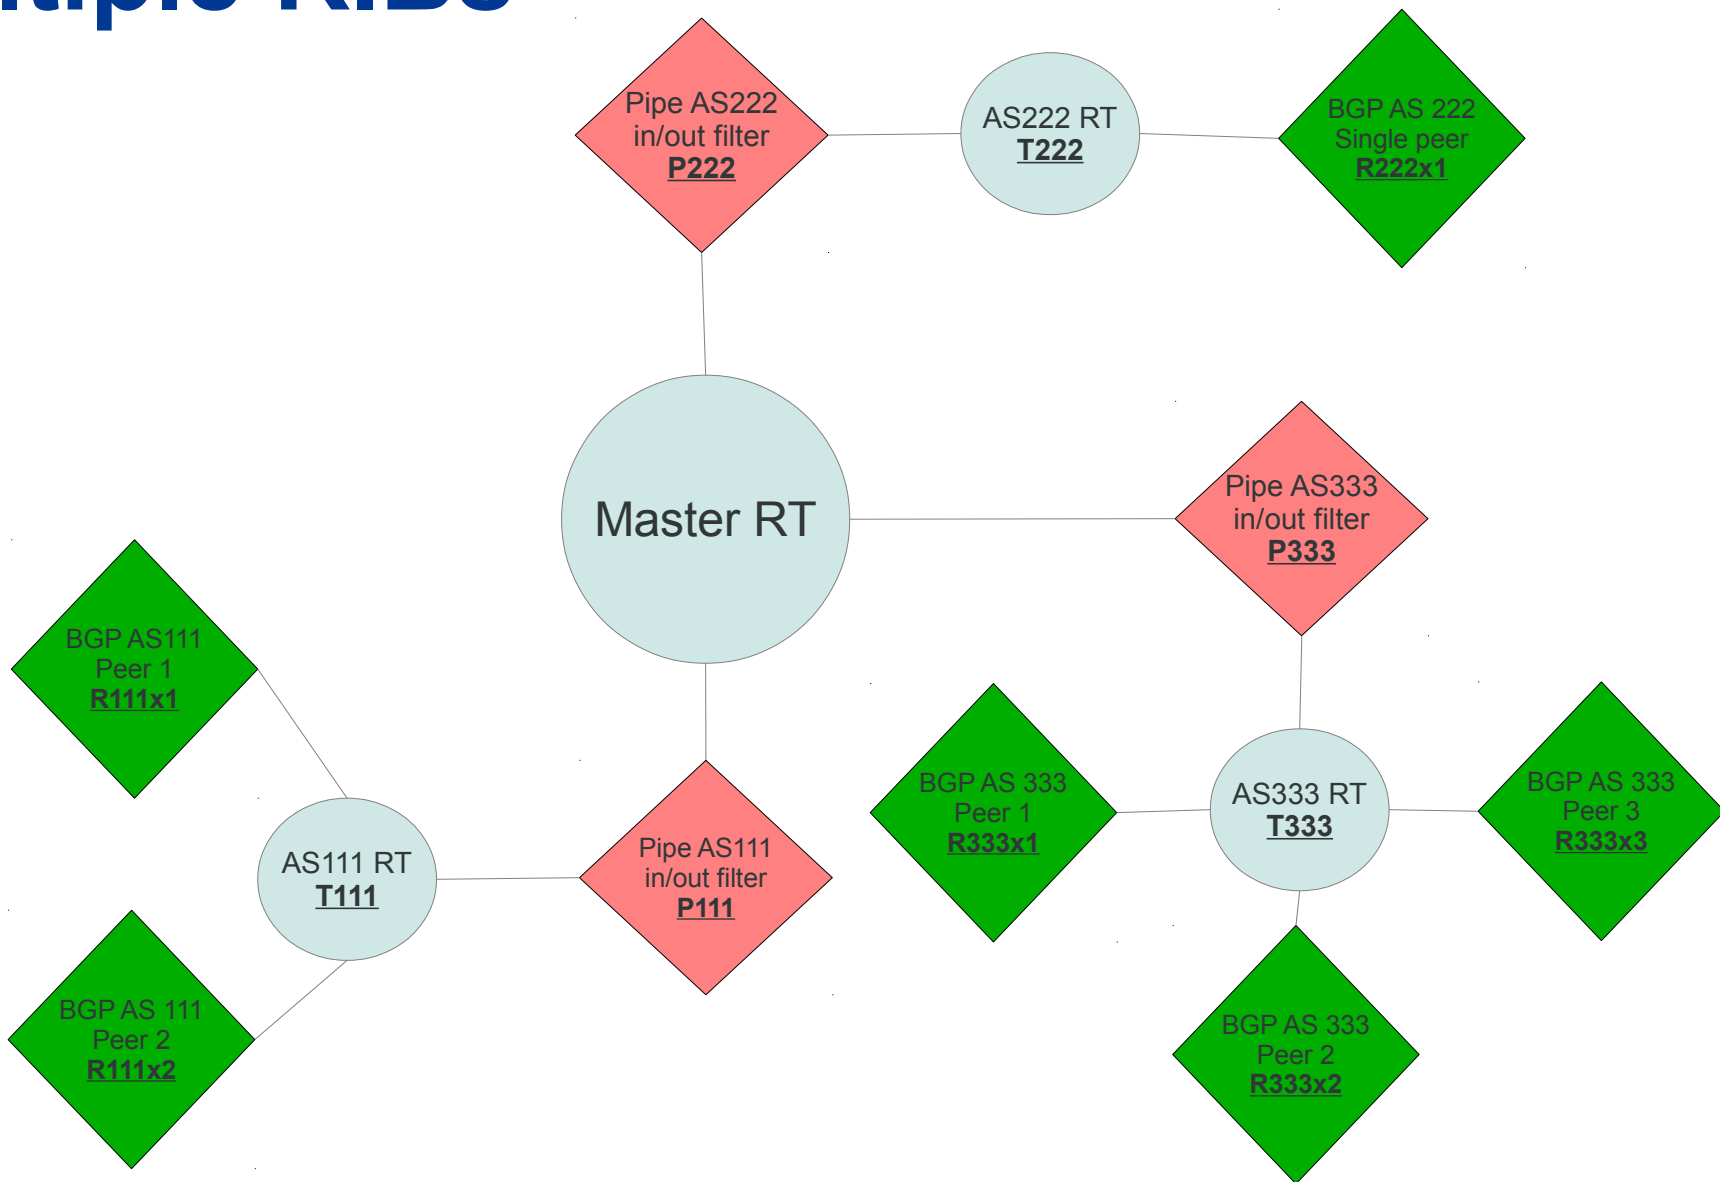


# BTW: IXP - Route filtering problem

- Normally filters match best route only
- If best route is denied no other is 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

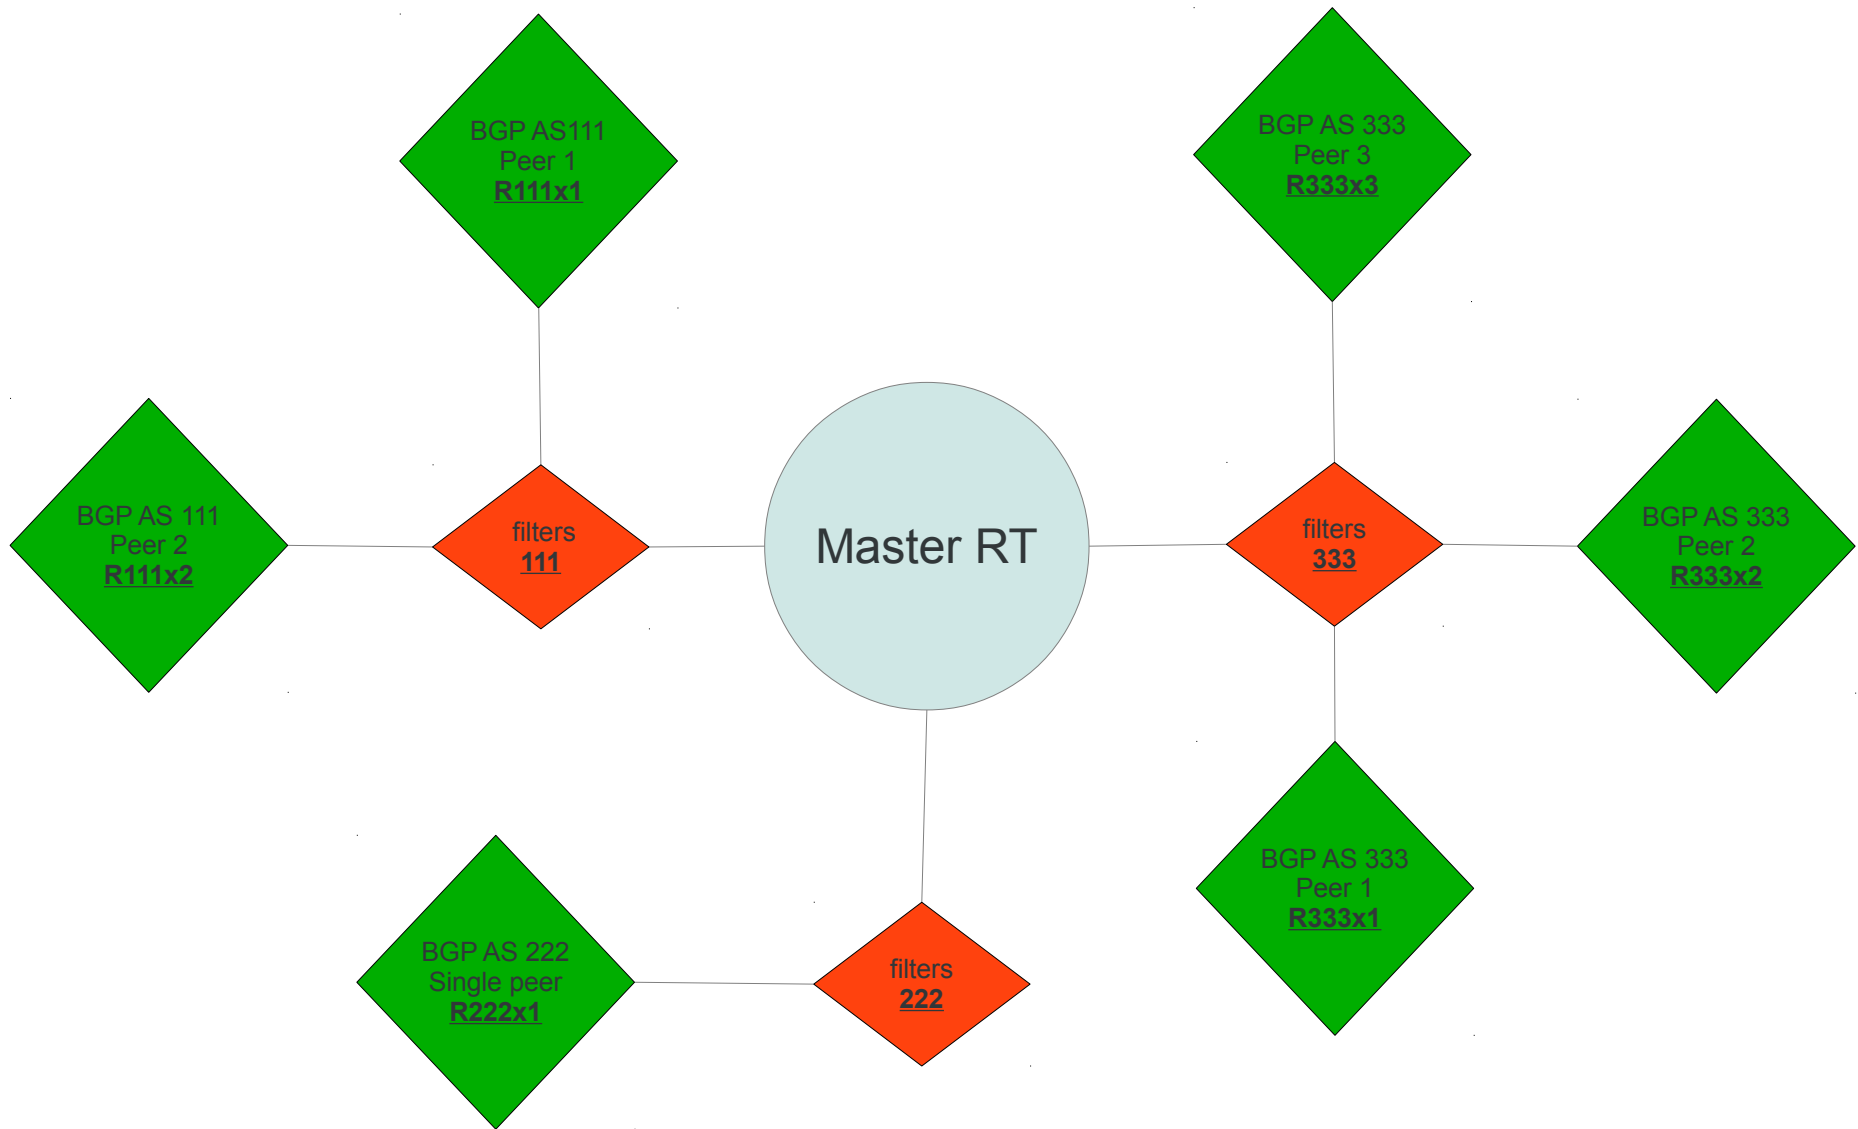


# Single RIB

- Allow second best route to be exported
- BGP protocol option **secondary**;
- Precondition: routing table sorted (a little bit slower)
  - **table master sorted**;
- No PIPE protocol
- To see filtered routes (1.3.9)
  - import keep filtered
  - show route filtered



# Single RIB





```
template bgp NIXPEERS {
    local as myas;
    rs client;
    secondary;
    import limit 25000 action disable;
    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 (~120 peers)
- ~7.000 IPv6 routes (~100 peers)
- Memory consumption 1/7 IPv4 and 1/4 IPv6



# Work in Progress



- Lightweight client (less libraries) - OpenWRT
- BGP Add-path (beta)
- Universal looking glass
- IPv4 and IPv6 integration – 2.0.0
- IS-IS (after integration)
  
- ... your feedback is VERY welcome

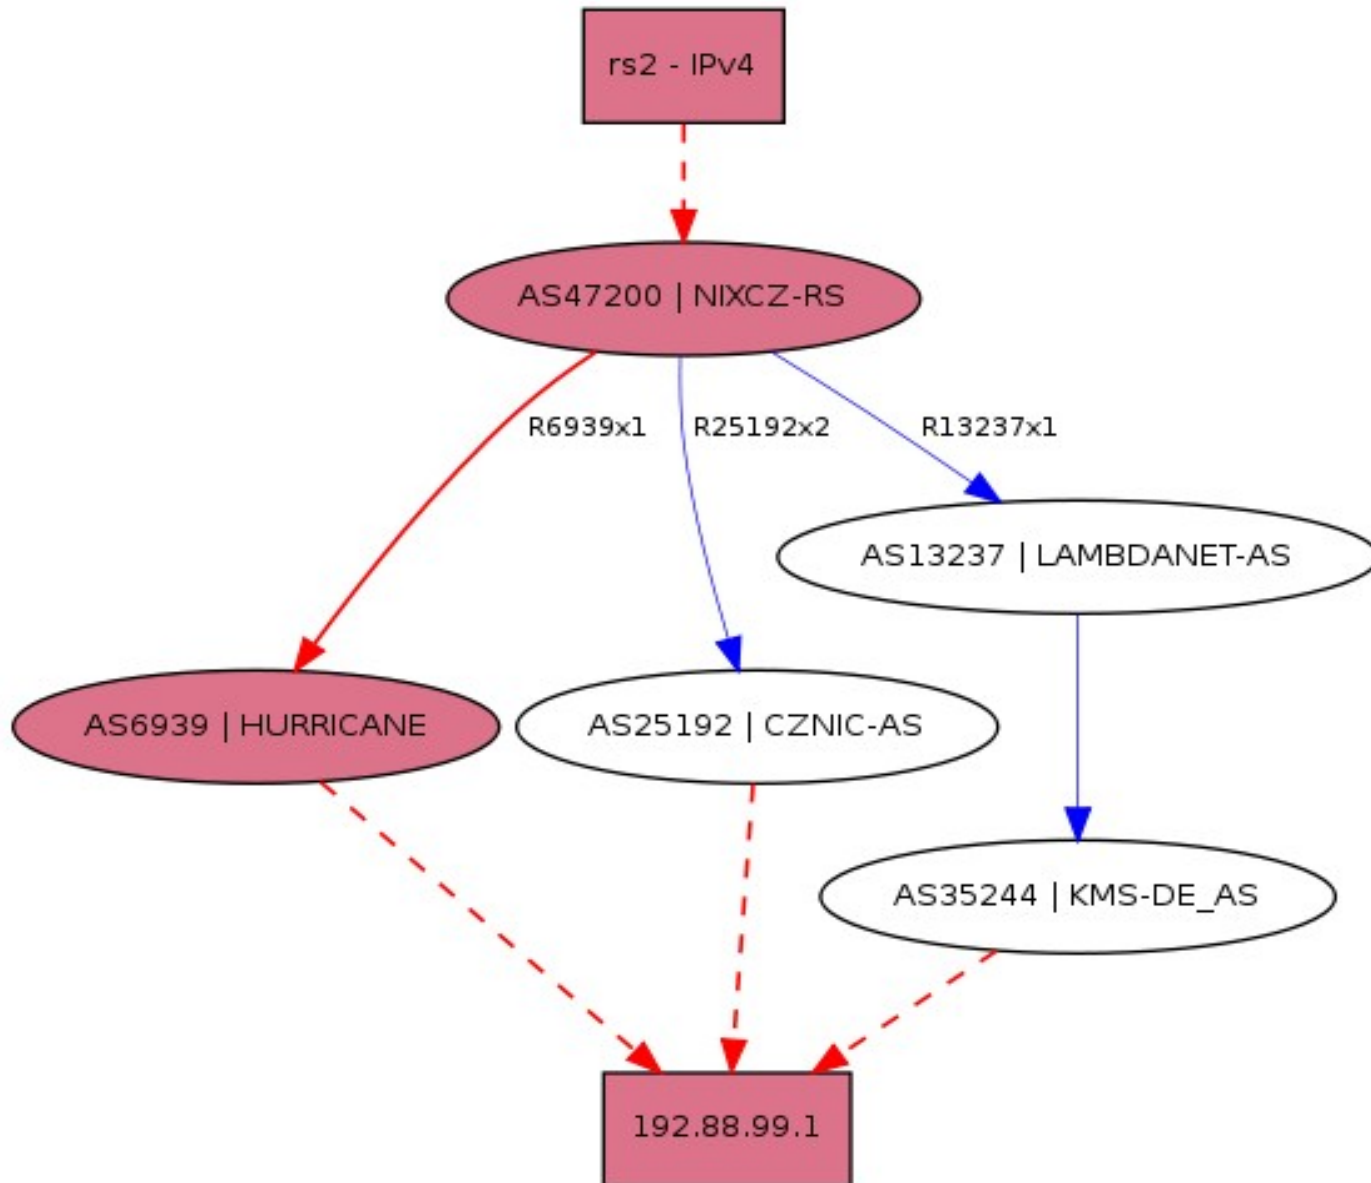


# Universal Looking Glass

- Supports multiple routers
  - BIRD
  - Cisco
  - Juniper
  - ... more to come
- Path visualisation



# Universal Looking Glass



# Conclusion



- BIRD stable and widely deployed
- Many new BIRD features
- Secondary route export can save resources
  - More important for larger deployments
- And look forward to more :-)
- Check wiki at <http://bird.network.cz>
- Feedback welcome!





**Thank You!**



**Ondrej Filip • [ondrej.filip@nic.cz](mailto:ondrej.filip@nic.cz) • <http://bird.network.cz>**

