

# BIRD's flight from Lisbon to Prague



CZ.NIC

Ondrej Filip / [ondrej.filip@nic.cz](mailto:ondrej.filip@nic.cz)

With help of Wolfgang Hennerbichler


May 6, 2009 – RIPE 60 / Prague

# Features



- GPL implementation of RIP, BGP, OSPF
- BGP – ASN32, MD5, route server, ...
- IPv4, IPv6
- Highly flexible – Multiple routing tables - RIBs (internal and also synchronization with OS)
- Protocol PIPE - multiple routers, route reflectors
- Powerful configuration & filtering language
- Automatic reconfiguration
- Currently developed by CZ.NIC Labs

# Design



# Filters example



```
define myas = 47200;
```

```
function bgp_out(int peeras)
{
  if ! (source = RTS_BGP ) then return false;
  if (0,peeras) ~ bgp_community then return false;
  if (myas,peeras) ~ bgp_community then return true;
  if (0, myas) ~ bgp_community then return false;
  return true;
}
```

```
protocol bgp R25192x1 {
  local as myas;
  neighbor 194.50.100.13 as 25192;
  import where bgp_in(25192);
  export where bgp_out(25192);
  rs client;
}
```

# Filters example

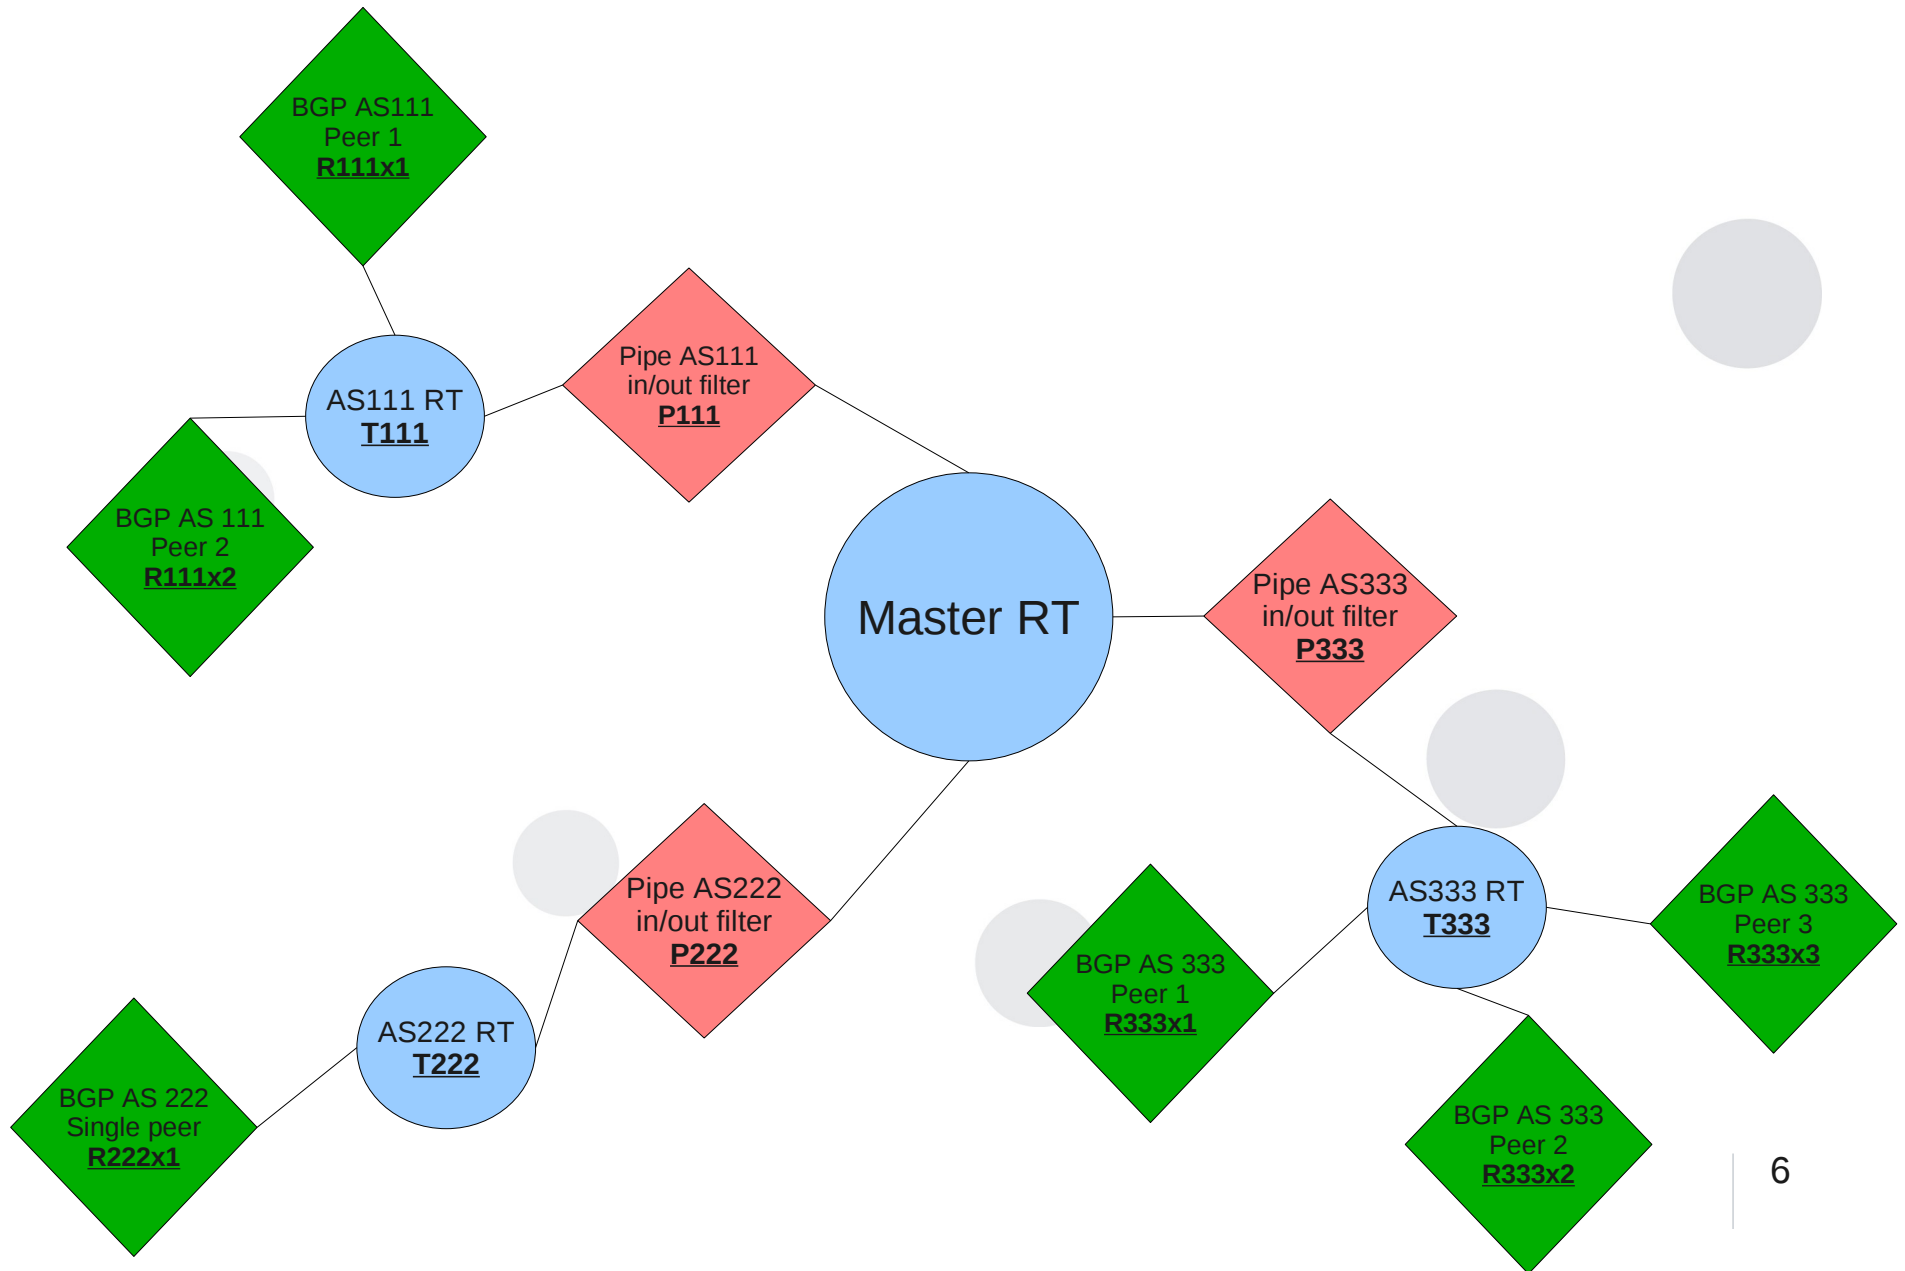


```
function avoid_martians()
prefix set martians;
{
    martians = [ 169.254.0.0/16+, 172.16.0.0/12+,
        192.168.0.0/16+, 10.0.0.0/8+, 224.0.0.0/4+,
        240.0.0.0/4+, 0.0.0.0/32-, 0.0.0.0/0{25,32},
        0.0.0.0/0{0,7} ];

    # Avoid RFC1918 networks
    if net ~ martians then return false;

    return true;
}
```

# BIRD as a route server



# BIRD at LoNAP

- First BIRD RS implementation
- Two route-servers
- BIRD and OpenBGPD
- BIRD in multiple RIBs setup
- Thanks for help with debugging!
- 27 sessions, 2500 prefixes

# BIRD at NIX.CZ



- Also multiple RIBs setup
- BIRD on Linux and Quagga on FreeBSD
- Approx 100 IPv4 sessions (30 IPv6)
- No BIRD's crash since implementation (1.1.3 release)
- Approx 8000 IPv4 prefixes (120 IPv6)
- Memory consumption – 80MB



# BIRD at LINX



- Some bugfixes sent to the developers – thank you!
- Use SNMP Scripts to monitor routes
- BGP communities
- 190 BGP sessions, 50k prefixes

# BIRD at MSK-IX



- Reported bugs to the BIRD team (crashes related to bug in Cisco BGP) fixed
- Some additional feature request – time format, no-export communities etc.
- Deals with ASN32 customers (also has a translation table for ASN32 communities)
- 299 members, 11800 prefixes, ~600MB BIRD memory footprint, IPv4 only
- Implementation of own **Looking Glass** – 'show bgp summary'

# BIRD at DECIX



- Stable BIRD instance
- BIRD running on Debian Lenny
- BIRD daemon uses 3.4G of RAM at 380 BGP sessions with multiple RIBs
- Controlled via BGP Communities
- AS32 supported, IPv6 supported

# BIRD at VIX



- Stable, but not fully released yet (test-phase with 10 members), to be released soon
- Fully configurable through webinterface (no BGP communities)
- Debian on XEN (timing issues fixed)
- Fully IPv6 aware
- Prefixes from RS are always site local, can be prepended via webinterface

# BIRD at PAIX

- No communities
- Single RIB
- IRR prefix filtering
- Multiple sites with central configuration systems
- (PAO, SEA, CHI, DAL, NYC, ATL, TOR, TYS)

# BIRD at other IXPs

- Implemented
  - MINAP
  - ECIX
- Tested by JPNAP and some others...

# New features since 09/09/09

- OSPFv3
- BGP passive option
- Improved filtering (mainly BGP related)
- Read only CLI
- Richer 'show' command
- RFC5004 (prefer older external routes)
- RFC2918 (Route Refresh)
- MRTdumping – BGP only

**59 -> 60**

# Future

- New web pages design
- Examples in documentation
- IPv6 Route advertisement
- Lightweight CLI (OpenWRT)
- Route flap dampening
- Solaris (CZ.NIC DNS anycasting)
- ....
- Depends mainly on users request



# Conclusion (for IXPs)

- BIRD is stable as a route server
- Deployed by many IXPs
- Also deployed by biggest IXPs
- Flexible configuration for your needs
- Send us feedback of feature request!

# ¿Questions?

<http://labs.nic.cz>  
<http://bird.network.cz>  
<[ondrej.filip@nic.cz](mailto:ondrej.filip@nic.cz)>