



# BIRD Internet Routing Daemon

Route Servers



Ondrej Filip • [ondrej.filip@nic.cz](mailto:ondrej.filip@nic.cz)  
23 May 2013 • CEE Peering Day • Prague

# CZ.NIC, CZ.NIC Labs



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



# BIRD Project goals (1999)



- Opensource routing daemon – alternative to Quagga/Zebra (and GateD at that time)
- Fast and efficient
- Portable, modular
- Support current routing protocols
- IPv6 and IPv4 in one source code – dual compilation
- Easy reconfiguration and powerful filtering





# Features

- Portable – Linux, FreeBSD, NetBSD, OpenBSD
- IPv4/IPv6 support, Static routing
- RIP, RIPv2, RIPng
- OSPFv2, OSPFv3
- MRTdump logging
- IPv6 router advertisements
- BGP, BGP route server, BGP route reflector
- 4B ASNs, BGP extended community support
- ROA support (RPKI)



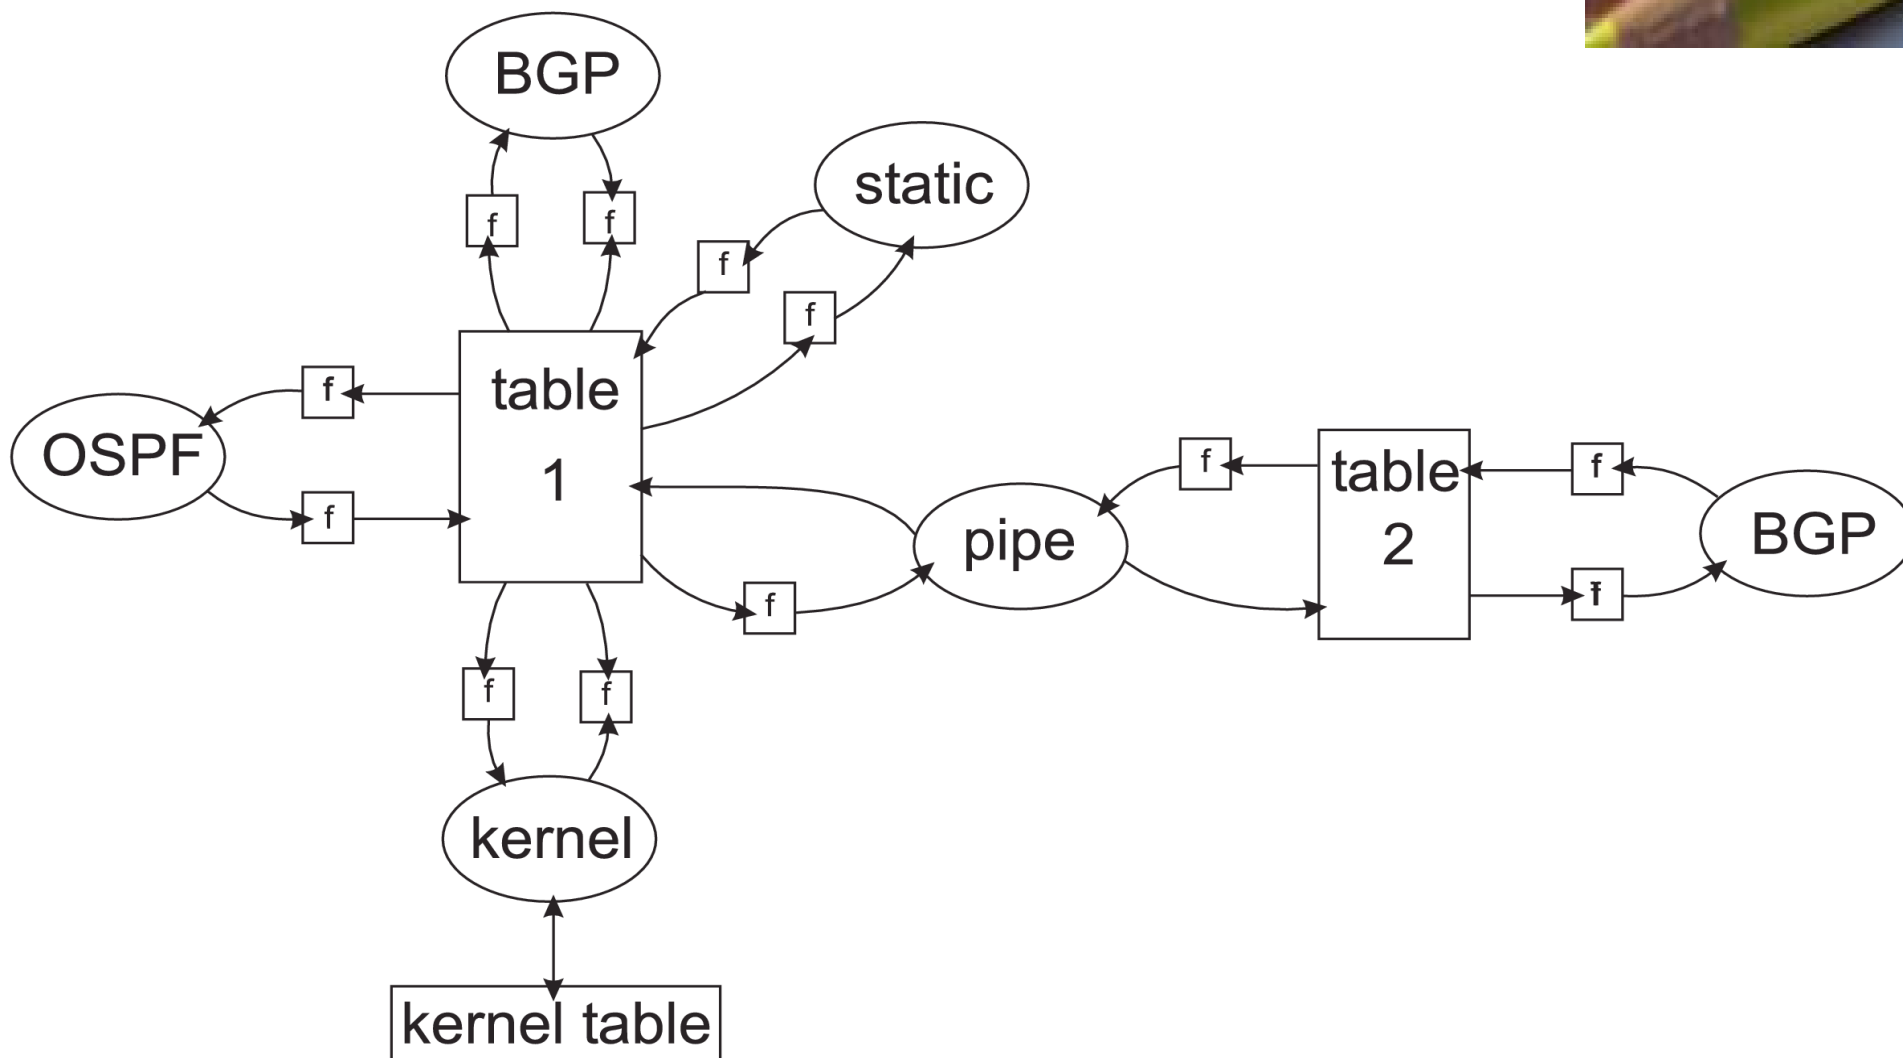


# Features

- Multiple routing table - RIBs (internal and also synchronization with OS)
- Protocol PIPE
- Multiple routers, route reflectors on a single system
- Powerful configuration
  - C-language style, curly brackets
- Very powerful filtering language
- Command line interface (show, restart, ...)
- Smart reconfiguration



# Design



# Configuration sample

```
log "/var/log/bird.log" all;

router id 193.51.100.238;

protocol static {
    route 10.0.0.0/8 drop;
    route 172.16.0.0/12 drop;
    route 192.168.0.0/16 drop;
}
filter bgp_out {
    if (net = 192.175.48.0/24 ) &&
        (source = RTS_DEVICE) then accept;
    else reject;
}
protocol bgp NIX_1 {
    local as 112;
    neighbor 193.51.100.235 as 6981;
    import all;
    export filter bgp_out;
}
```



# CLI example



```
bird> show protocols
```

name	proto	table	state	since	info
direct1	Direct	master	up	Apr11	
kernell	Kernel	master	up	Apr11	
devicel	Device	master	up	Apr11	
static1	Static	master	up	Apr11	
NIX_2	BGP	master	up	Apr11	Established
NIX_1	BGP	master	up	Apr25	Established
ospf1	OSPF	master	up	Apr11	Running

```
bird>
```

```
bird> show status
```

```
BIRD 1.1.3
```

```
Current server time is 06-08-2009 22:01:06
```

```
Last reboot on 11-07-2009 22:54:12
```

```
Last reconfiguration on 30-07-2009 06:25:25
```

```
Daemon is up and running
```

```
bird>
```





# CLI example



```
bird> show route for 127.0.0.1
127.0.0.0/8          dev lo [direct1 13:09] * (240)

bird> show route filter bgp_out
192.175.48.0/24     dev dummy0 [direct1 Apr1] * (240)

bird> show route count
1469 of 1469 routes for 849 networks

bird> show route export NIX_1
192.175.48.0/24     dev dummy0 [direct1 Apr1] * (240)

bird> show route where 127.0.0.5 ~ net
0.0.0.0/0           via 195.47.235.1 on eth0 [static Apr1] * (200)
127.0.0.0/8        dev lo [direct1 Apr1] * (240)

bird> show route filter {if 127.0.0.5 ~ net then accept;}
0.0.0.0/0           via 195.47.235.1 on eth0 [static Apr1] * (200)
127.0.0.0/8        dev lo [direct1 Apr1] * (240)
```



# Filter 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 and similar networks  
    if net ~ martians then return false;  
  
    return true;  
}
```




# Filters




- Filters compiled into byte code
- Variables!, Sets
- Data types: bool, int, ip, prefix, enum, quad, string, bgppath, bgpmask, clist
- Operators: +, -, \*, /, comparisons, logical, element\_of\_set (~)
- Control structures – if/else, case
- Complicated structures compiled into balanced binary trees (or similar) – log time complexity



# IXP without route server



# IXP with route server



# Route servers

- Without RS
  - Each IXP member –  $n-1$  sessions = total  $n*(n-1)/2$
  - Setup with every new member
  - Main router CPU load (esp. with TCP-MD5)
- With RS
  - Just 1 BGP session per member router (or possibly 2)
  - Decreases the number of BGP sessions
  - IXP attractive for newcomers!
  - Very strong and complex **filtering**
  - Not loosing routing policy – BGP community signalling



# Route server policy

```
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;  
}
```



# BIRD at NIX.CZ



- AS112
- Dual route server
- Linux & FreeBSD, IPv4 & IPv6
- About 120 BGP sessions
- Configuration taken from intranet and RPSL (RIPE DB)
- Filtering based on prefix, AS path pairs
- Reconfigured every 2 hours
- More than 6M characters





# Deployed at ... (and more)



# BIRD - Work in Progress



- BGP Add-path (beta)
- IPv4 and IPv6 integration – 2.0.0
- IS-IS (after integration)
  
- ... your feedback is VERY welcome

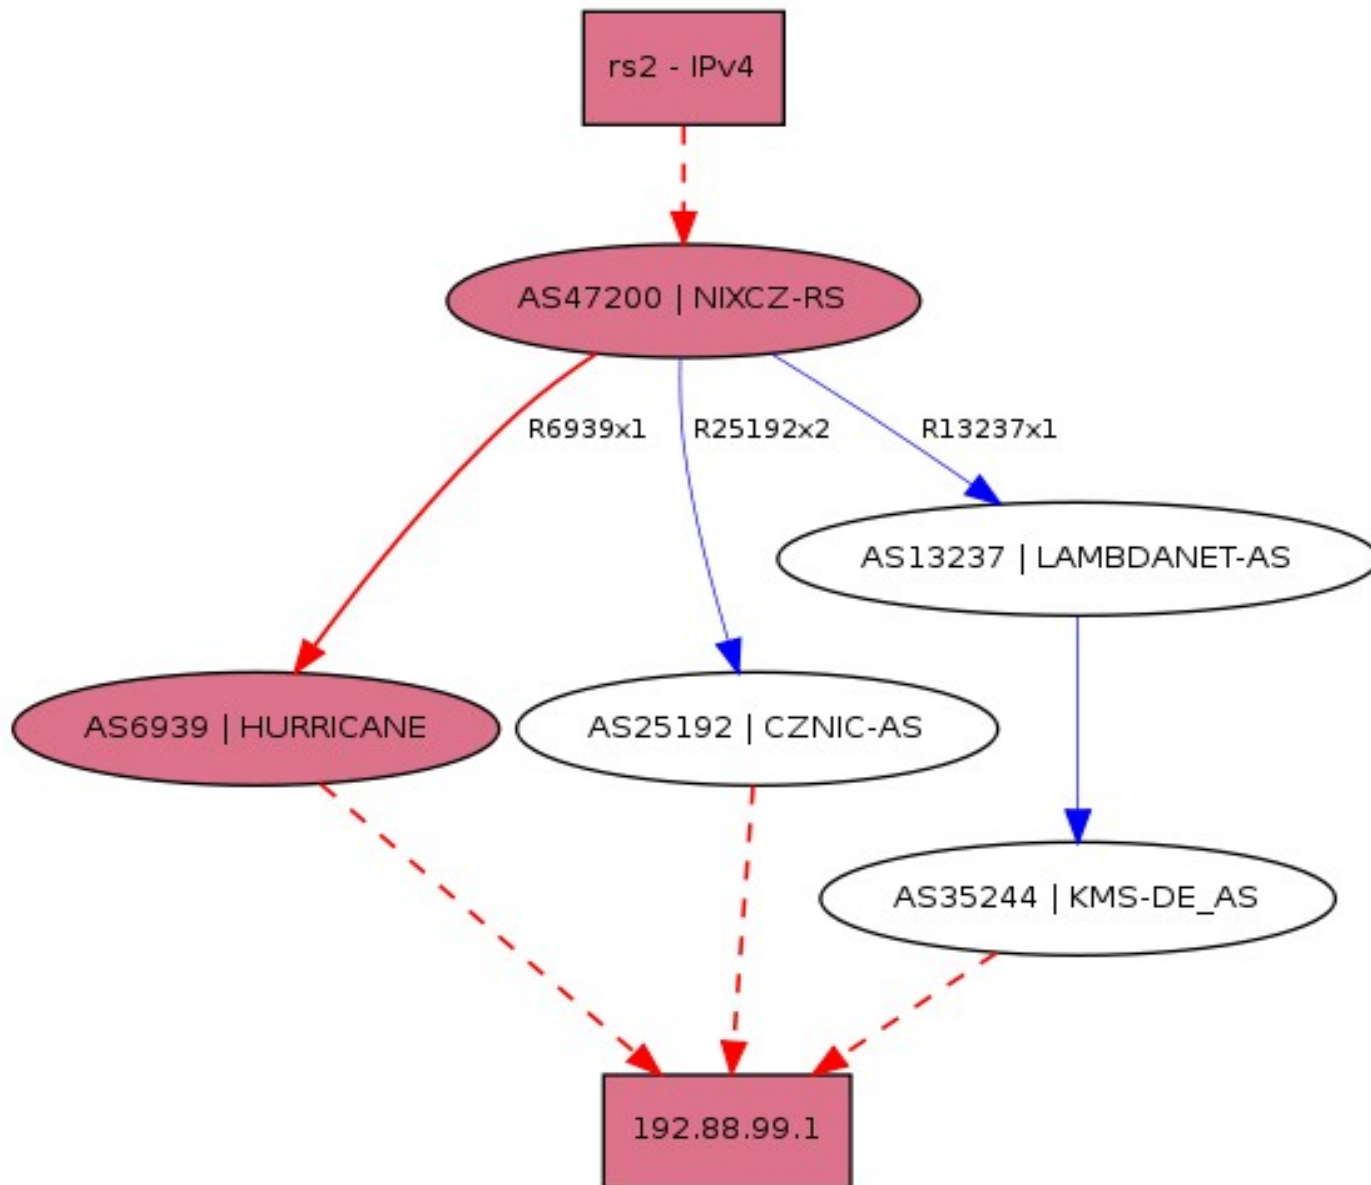


# Universal Looking Glass

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



# Universal Looking Glass



# Conclusion



- BIRD stable and widely deployed
- BIRD a route server – developed by team with IXP knowledge
- Filtering at RS level
- RS make IXP attractive for newcomers
- Check <http://bird.network.cz>
  
- Feedback welcome!





**Thank You!**



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

