



# Extended communities on RS

Ondrej Filip / of@nix.cz  
NIX.CZ

Nov 3, 2011, RIPE, EIX, Wien


# NIX.CZ is 15!

- ➔ You are welcome!!!
- ➔ November 21, 2011
- ➔ Planetarium, historic tram, party, ...
- ➔ Contact us if you're interested to come

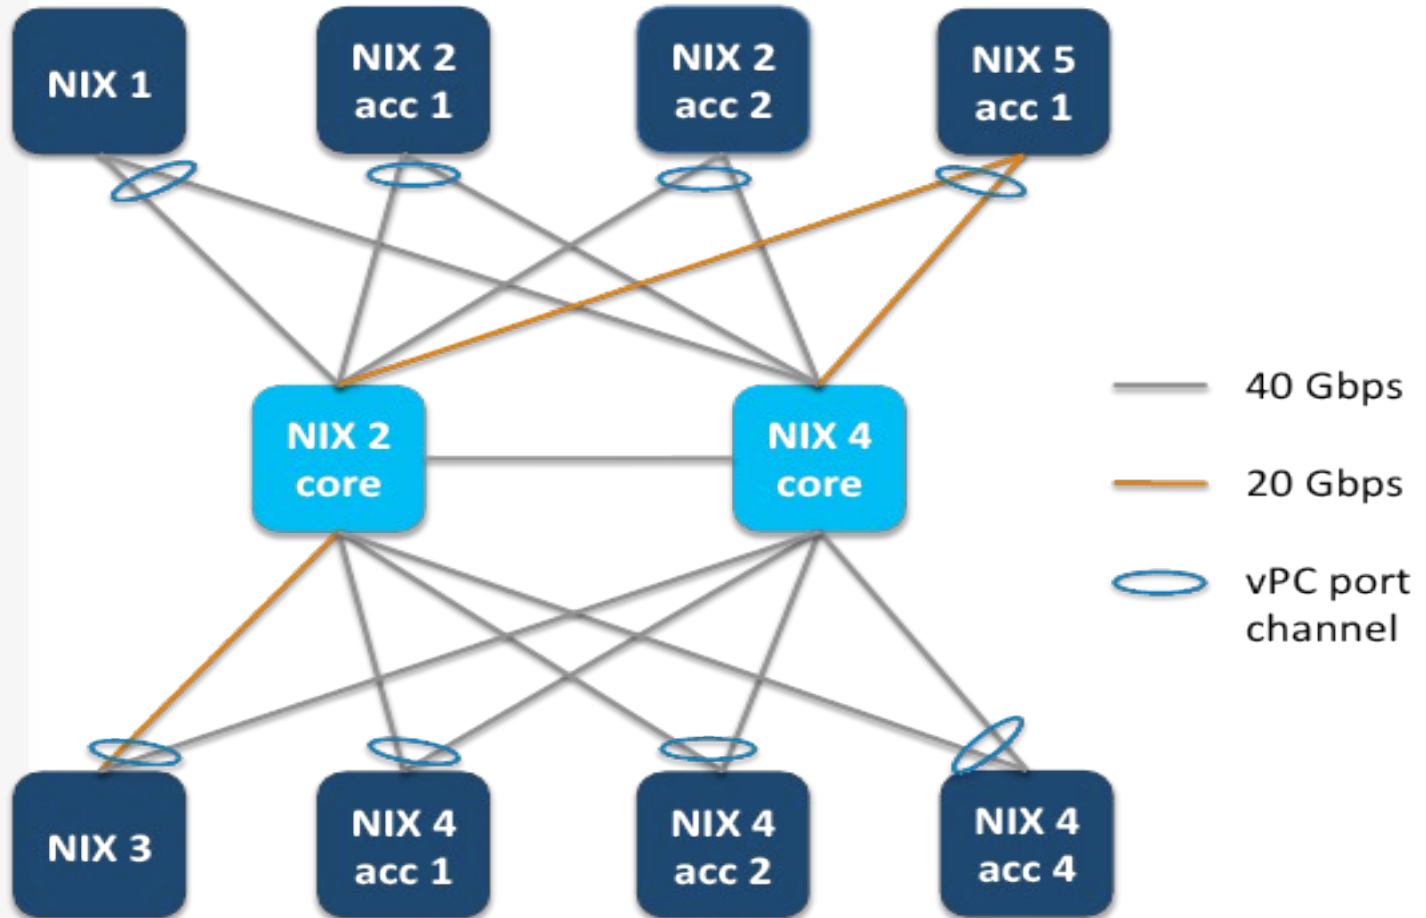


# NIX.CZ update

- ➔ 200 Gbps reached!
- ➔ A lot of new customers/ports
- ➔ New POP opened - 5 POPs
- ➔ Price list simplification and price reduction plan for 2012
- ➔ Switch upgrades - Cisco Nexus also used as an access switch



# NIX.CZ update



# ASN32



- ➔ Current RS policy based on standard BGP communities
- ➔ 16 bits only
- ➔ New members with ASN32
  - ⑤ They can filter others
  - ⑤ But they cannot be filtered

# Current RS policy

Evaluation order	Community	Action
1	0:<peer-as>	Do not advertise to <peer-as>
2	47200:<peer-as>	Advertise to <peer-as>
3	0:47200	Do not advertise to any peer
4	47200:47200	Advertise to all peers

# Extended community

- ➔ RFC 4360
- ➔ 48 bits for ASNs (16+32)
- ➔ Plus some signalling bits
- ➔ More types – route origin/route target
- ➔ AS number or IP address

# New RS policy

Eval. order	Community	Action
1	0:<peer-as> or So:0:<peer-as>	Do not advertise to <peer-as>
2	47200:<peer-as> or So:47200:<peer-as>	Advertise to <peer-as>
3	0:47200 or So:0:47200	Do not advertise to any peer
4	47200:47200 or So:47200:47200	Advertise to all peers



# New RS policy

- ➔ Conditions evaluated in fixed order (1-4)
- ➔ You can mix both – not recommended
- ➔ 'OR' clause applied
- ➔ All (extended) communities containing '47200:' or '0:' removed by route server

# So:0.0.0.0:<peer-as>

- ➔ Some Cisco CRS boxes cannot generate So:0:<ASN32>
- ➔ RFC not very clear on that
- ➔ Those members propagate So:0.0.0.0:<ASN16> instead
- ➔ Cannot 'block' any peer
- ➔ Peers can be only included
- ➔ Example:
  - ⑤ So:47200:peer1
  - So:47200:peer2
  - So:47200:peer3
  - So:0.0.0.0:47200

# NIX.CZ Route Servers

- ➔ Two route servers
- ➔ BIRD & Quagga – implemented on both
- ➔ Problems with Quagga – communities removal, reconfiguration, filtering etc.

# BIRD

- ➔ Extended communities introduced in 1.3.3
- ➔ New release 1.3.4 - bugfix release - recommended
- ➔ Tested and run at NIX.CZ
- ➔ RS configuration templates to be published soon
- ➔ Other new features - 'include', BGP TTL security, OSPF NSSA, static recursive routes



# Conclusion

- ➔ Extended communities used for RS policy at NIX.CZ
- ➔ ASN32 fully supported
- ➔ Problems with Cisco - reported - we can change the policy
- ➔ BIRD works
- ➔ Quagga sort of works :-)
- ➔ OpenBPGd untested yes



# Questions ?

<http://www.nix.cz>

<http://bird.network.cz>

<mailto:of@nix.cz>

