

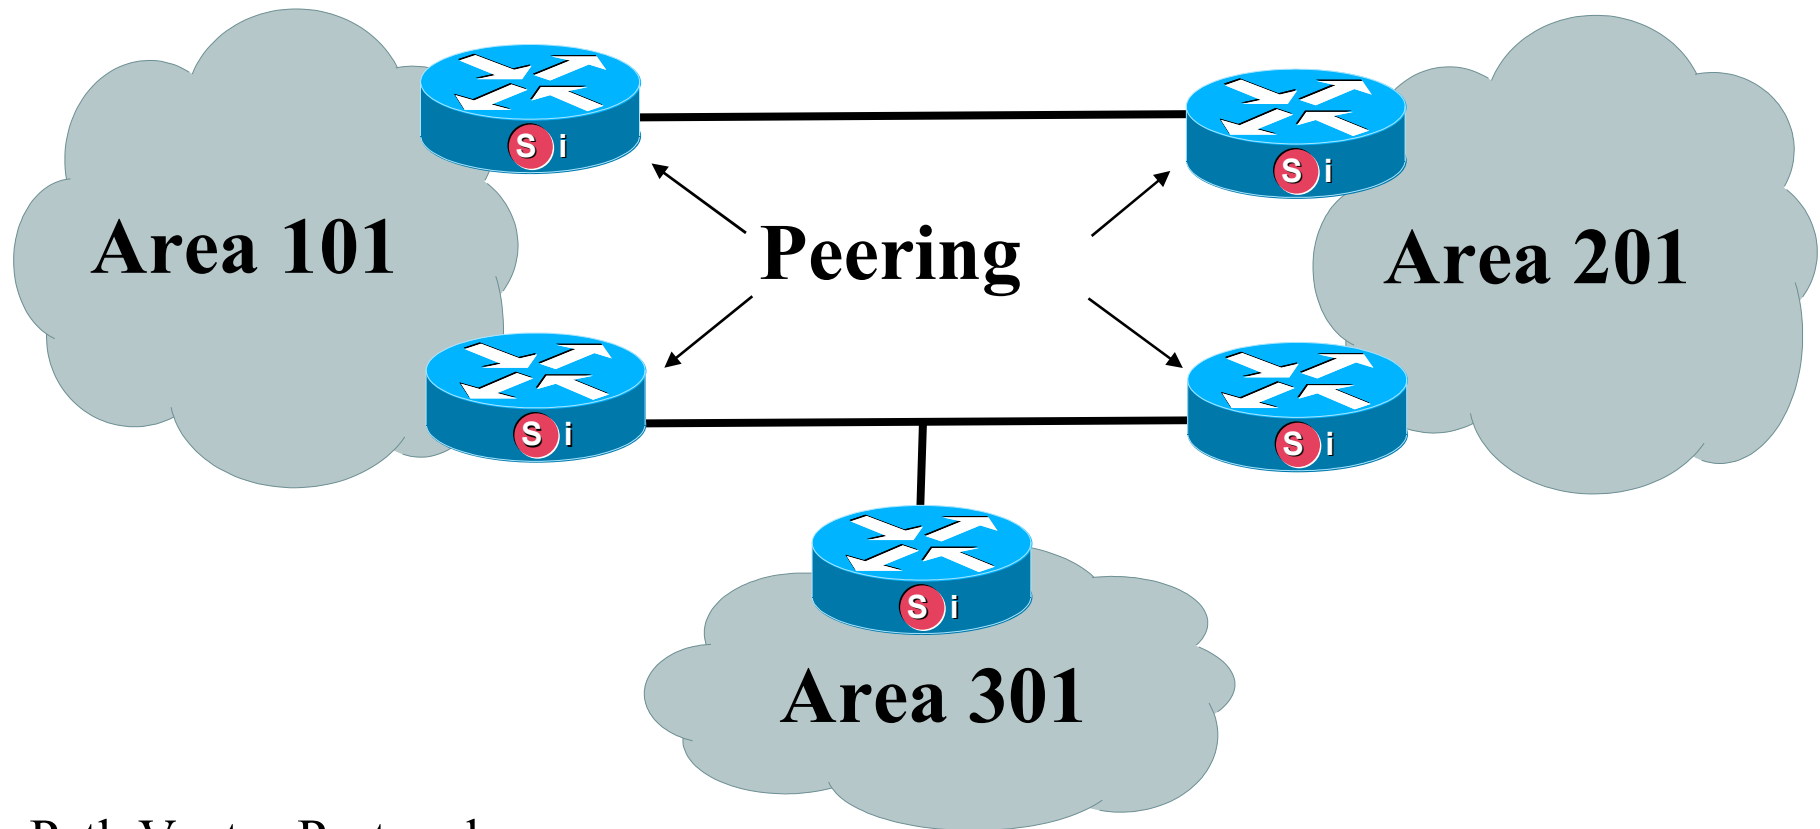
IPv6 Routing: BGP

János Mohácsi
NIIF/HUNGARNET

Outline

- Introduction to MBGP
- Configuration examples
- Route filtering
- RPSLng

Multiprotocol BGP - Overview



- Path Vector Protocol
 - Carries sequence of AS numbers indicating path
- Ties Autonomous Systems together via Peering
- Multiple address families: ipv4, ipv6, unicast, multicast, vpns etc.

BGP-4 Extensions for IPv6 (RFC 2545)

- BGP-4 carries only 3 pieces of information which is truly IPv4 specific:
 - NLRI in the UPDATE message contains an IPv4 prefix
 - NEXT_HOP path attribute in the UPDATE message contains a IPv4 address
 - BGP Identifier is in the OPEN message & AGGREGATOR attribute
- To make BGP-4 available for other network layer protocols, RFC 2858 (obsoletes RFC 2283) defines multi-protocol extensions for BGP-4
 - Enables BGP-4 to carry information of other protocols e.g. MPLS,IPv6
 - New BGP-4 optional and non-transitive attributes:
 - MP_REACH_NLRI
 - MP_UNREACH_NLRI
 - Protocol independent NEXT_HOP attribute
 - Protocol independent NLRI attribute

BGP-4 Extensions for IPv6

- New optional and non-transitive BGP attributes:

MP_REACH_NLRI (Attribute code: 14)

“Carry the set of reachable destinations together with the next-hop information to be used for forwarding to these destinations” (RFC2858)

MP_UNREACH_NLRI (Attribute code: 15)

Carry the set of unreachable destinations

- Attribute 14 and 15 contains one or more Triples:

Address Family Information (AFI)

Next-Hop Information (must be of the same address family)

NLRI

BGP-4 Extensions for IPv6

- Address Family Information (AFI) for IPv6

AFI = 2 (RFC 1700)

Sub-AFI = 1 Unicast

Sub-AFI = 2 (Multicast for RPF check)


Sub-AFI = 3 for both Unicast and Multicast

Sub-AFI = 4 Label

Sub-AFI = 128 VPN

BGP-4 Extensions for IPv6

- Next-hop contains a global IPv6 address or potentially a link local (for iBGP update this has to be changed to global IPv6 address with route-map)
- The value of the length of the next hop field on MP_REACH_NLRI attribute is set to 16 when only global is present and is set to 32 if link local is present as well
- Link local address as a next-hop is only set if the BGP peer shares the subnet with both routers (advertising and advertised)



BGP-4 Extensions for IPv6

- TCP Interaction

BGP-4 runs on top of TCP

This connection could be setup either over IPv4 or IPv6

- Router ID


When no IPv4 is configured, an explicit bgp router-id needs to be configured

This is needed as a BGP Identifier, this is used as a tie breaker, and is send within the OPEN message

Cisco BGP-4 Configurations for IPv6

Router A

```
router bgp 1
no bgp default ipv4 unicast
bgp router-id 1.1.1.1
neighbor 3ffe:b00:ffff:2::2 remote-as 2
address-family ipv6
neighbor 3ffe:b00:ffff:2::2 activate
network 2003:3:2::/64
network 2003:3:3::/64
```




Juniper BGP-4 Configurations for IPv6

Router A

```

routing-options {
  autonomous-system 1
  router-id 1.1.1.1;
}
protocols {
  bgp {
    group to_as2 {
      family inet6 {
        (any | multicast | unicast)
      }
      type external;
      peer-as 2;
      neighbor 3ffe:b00:ffff:2::2;
    }
  }
}

```



Policy routing with BGP

- Remove bogon routes:
 - IPv6 specific rules
 - 3FFE::/16 (6bone) treat according to the 6bone rules
 - 2002::/16 only permits the /16, no more-specifics
 - 0000::/8 is denied (loopback, unspecified, v4-mapped)
 - FE00::/9 and FF00::/8 are denied (multicast ranges, RFC3513)
 - all the rest of the IPv6 unicast address space (0::0/0) is permitted up to reasonable sizes

Cisco Prefix list for filtering

- Deny document prefix: 2001:db8::/32. Deny based on rules on previous slide . Accept more specific routes up to a /48 from RIR (2000::/3) ranges.

```
ipv6 prefix-list ipv6-ebgp-liberal deny 2001:db8::/32 le 128
ipv6 prefix-list ipv6-ebgp-liberal permit 2002::/16
ipv6 prefix-list ipv6-ebgp-liberal deny 2002::/16 le 128
ipv6 prefix-list ipv6-ebgp-liberal deny 0000::/8 le 128
ipv6 prefix-list ipv6-ebgp-liberal deny fe00::/9 le 128
ipv6 prefix-list ipv6-ebgp-liberal deny ff00::/8 le 128
ipv6 prefix-list ipv6-ebgp-liberal permit 0::/0 le 48
ipv6 prefix-list ipv6-ebgp-liberal deny 0::/0 le 128
```

Juniper policy for filtering

```
policy-statement ipv6-ebgp-liberal {  
  from {  
    family inet6;  
    route-filter ::/8 orlonger reject;  
    route-filter 2001:db8::/32 orlonger reject;  
    route-filter 2002::/16 longer reject;  
    route-filter fe00::/9 orlonger reject;  
    route-filter ff00::/8 orlonger reject;  
    route-filter ::/0 upto /48 next policy;  
  }  
  then reject;  
}
```

RPSLng

- Need for operate the network properly: Specifying and make publicly available the routing policy
 - Who are the peers
 - What routes are
 - Originated by peer
 - Imported
 - Exported
 - Preferred when multiple routes exist
 - What routes to aggregate
 - What to do if no route exists

Route Object

- Use CIDR length format
- Specifies origin AS for a route
- Can indicate membership of a route set

route:	[mandatory]	[single]	[primary/look-up key]
descr:	[mandatory]	[multiple]	
origin:	[mandatory]	[single]	[primary/inverse key]
withdrawn:	[optional]	[single]	
member-of:	[optional]	[single]	[inverse key]
inject:	[optional]	[multiple]	
components:	[optional]	[single]	
aggr-bndry:	[optional]	[single]	[inverse key]
aggr-mtd:	[optional]	[single]	
export-comps:	[optional]	[single]	
holes:	[optional]	[single]	
remarks:	[optional]	[multiple]	
cross-nfy:	[optional]	[multiple]	[inverse key]
cross-mnt:	[optional]	[multiple]	[inverse key]
notify:	[optional]	[multiple]	[inverse key]
mnt-by:	[mandatory]	[multiple]	[inverse key]
changed:	[mandatory]	[multiple]	
source:	[mandatory]	[single]	

AS Set

- Collect together Autonomous Systems with shared properties
- Can be used in policy in place of AS
- RPSL has hierarchical names

<code>as-set:</code>	<code>[mandatory]</code>	<code>[single]</code>	<code>[primary/look-up key]</code>
<code>descr:</code>	<code>[mandatory]</code>	<code>[multiple]</code>	
<code>members:</code>	<code>[optional]</code>	<code>[single]</code>	
<code>mbrs-by-ref:</code>	<code>[optional]</code>	<code>[single]</code>	
<code>remarks:</code>	<code>[optional]</code>	<code>[multiple]</code>	
<code>tech-c:</code>	<code>[mandatory]</code>	<code>[multiple]</code>	<code>[inverse key]</code>
<code>admin-c:</code>	<code>[mandatory]</code>	<code>[multiple]</code>	<code>[inverse key]</code>
<code>notify:</code>	<code>[optional]</code>	<code>[multiple]</code>	<code>[inverse key]</code>
<code>mnt-by:</code>	<code>[mandatory]</code>	<code>[multiple]</code>	<code>[inverse key]</code>
<code>changed:</code>	<code>[mandatory]</code>	<code>[multiple]</code>	
<code>source:</code>	<code>[mandatory]</code>	<code>[single]</code>	

Route Set

- Collects routes together with similar properties

<code>route-set:</code>	<code>[mandatory]</code>	<code>[single]</code>	<code>[primary/look-up key]</code>
<code>descr:</code>	<code>[mandatory]</code>	<code>[multiple]</code>	
<code>members:</code>	<code>[optional]</code>	<code>[single]</code>	
<code>mbrs-by-ref:</code>	<code>[optional]</code>	<code>[single]</code>	
<code>remarks:</code>	<code>[optional]</code>	<code>[multiple]</code>	
<code>tech-c:</code>	<code>[mandatory]</code>	<code>[multiple]</code>	<code>[inverse key]</code>
<code>admin-c:</code>	<code>[mandatory]</code>	<code>[multiple]</code>	<code>[inverse key]</code>
<code>notify:</code>	<code>[optional]</code>	<code>[multiple]</code>	<code>[inverse key]</code>
<code>mnt-by:</code>	<code>[mandatory]</code>	<code>[multiple]</code>	<code>[inverse key]</code>
<code>changed:</code>	<code>[mandatory]</code>	<code>[multiple]</code>	
<code>source:</code>	<code>[mandatory]</code>	<code>[single]</code>	

Autonomous System Object

- Routing Policy Description object
- Most important components are
 - import
 - export
- Extension for multiprotocol bgp:
 - mp-import afi {ipv4/ipv6/any}. {unicast/multicast/any}
 - mp-export {ipv4/ipv6/any}. {unicast/multicast/any}
- These define the incoming and outgoing routing announcement relationships

Autonomous System Object (cont)

aut-num:	[mandatory]	[single]	[primary/look-up key]
as-name:	[mandatory]	[single]	
descr:	[mandatory]	[multiple]	
member-of:	[optional]	[single]	[inverse key]
import:	[optional]	[multiple]	[inverse key]
export:	[optional]	[multiple]	[inverse key]
default:	[optional]	[multiple]	[inverse key]
admin-c:	[mandatory]	[multiple]	[inverse key]
tech-c:	[mandatory]	[multiple]	[inverse key]
remarks:	[optional]	[multiple]	
cross-nfy:	[optional]	[multiple]	[inverse key]
cross-mnt:	[optional]	[multiple]	[inverse key]
notify:	[optional]	[multiple]	[inverse key]
mnt-by:	[mandatory]	[multiple]	[inverse key]
changed:	[mandatory]	[multiple]	
source:	[mandatory]	[single]	

Example RPSLng

```
aut-num: AS559
as-name: SWITCH
descr: SWITCH, Swiss Education and Research Network
mp-import: afi ipv6.unicast from AS3257 action pref=100; accept AS-TISCALI;
mp-import: afi ipv4.multicast from AS3303 action pref=100; accept AS-SWCMGLOBAL;
mp-import: afi ipv6.unicast from AS3303 action pref=100; accept AS-SWCMGLOBAL;
mp-import: afi ipv6.unicast from AS4589 action pref=100; accept AS-EASYNET;
mp-import: afi ipv6.unicast,ipv6.multicast from AS6680 action pref=106; accept ANY;
mp-import: afi ipv6.unicast from AS8758 action pref=100; accept AS-DOLPHINS;
mp-import: afi ipv6.unicast from AS9044 action pref=70; accept ANY;
mp-import: afi ipv4.multicast from AS20932 action pref=100; accept AS-IP-MAN;
mp-import: afi ipv4.multicast,ipv6.unicast from AS20965 action pref=106; accept AS-GEANTNRN AS-
GEANTEXTRAS AS-GEANTNRENPEERS;
mp-import: afi ipv4.multicast,ipv6.unicast from AS20965 action pref=100; accept AS-INFONET;
mp-export: afi ipv6.unicast to AS3257 announce AS-SWITCH;
mp-export: afi ipv4.multicast to AS3303 announce AS-SWITCH;
mp-export: afi ipv6.unicast to AS3303 announce AS-SWITCH;
mp-export: afi ipv6.unicast to AS4589 announce AS-SWITCH;
mp-export: afi ipv6.unicast,ipv6.multicast to AS6680 announce AS-SWITCH;
mp-export: afi ipv6.unicast to AS8758 announce AS-SWITCH;
mp-export: afi ipv6.unicast to AS9044 announce ANY;
mp-export: afi ipv4.multicast to AS20932 announce AS-SWITCH;
mp-export: afi ipv4.multicast,ipv6.unicast to AS20965 announce AS-SWITCH AS-CERNEXT;
default: to AS1299 action pref=10; networks ANY
default: to AS3549 action pref=10; networks ANY
admin-c: WH1101
tech-c: SNOCl-RIPE
mnt-by: SWITCH-MNT
mnt-lower: AS559-MNT
changed: simon@switch.ch 20040511
source: RIPE
```

Thank you!

- Acknowledgement to Simon Leinen and Gunter van de Velde for their help.
- Further informations:
 - <http://www.6net.org>
- Questions: mohacsi@niif.hu