Routing Protocols

Internal and External Routing

János Mohácsi (mohacsi@niif.hu)

Central Asia workshop,Ashgabat

Copy ... Rights

- This slide set is the ownership of the 6DISS project via its partners
- The Powerpoint version of this material may be reused and modified only with written authorization
- Using part of this material must mention 6DISS courtesy
- PDF files are available from www.6diss.org
- Looking for a contact ?
- Mail to : martin.potts@martel-consulting.ch
- Or bernard.tuy@renater.fr

Contributions

- Main authors
  - Carlos Friaças, FCCN, Portugal
  - Miguel Baptista, FCCN, Portugal
- Contributors
  - Mónica Domingues, FCCN, Portugal
  - Paulo Ferreira, FCCN, Portugal
  - János Mohácsi, NIIF/HUNGARNET, Hungary

Prerequisites

- You must have followed previously the modules:
  - 010-IPv6 Introduction
  - 020-IPv6 Protocol
  - 030-IPv6 Addressing

Agenda

- Internal Routing
  - Static Routing
  - RIPng
  - IS-IS
  - OSPFv3
- External Routing
  - Multiprotocol BGP

Static Routes

- Static route configuration syntax is the same as in IPv4
- Except Prefix and next-hop are IPv6

IPv4 static route:

ip route *[ipv4_prefix][ipv4_address_mask][ipv4_if_address]*

IPv6 static route:

ipv6 route *[ipv6_prefix/prefix_length][ipv6_if_address]*

ipv6 route ::/0 FastEthernet1/40 FE80::206:2AFF:FE58:7820

Static Routes

- It is not recommended to use a global unicast address as a next-hop addresss
- ICMPv6 redirect messages will not work if used

RFC 2461:

A router must be able to determine the link-local address of each of its neighboring routers in order to ensure that the target address of a Redirect message identifies the neighbor router by its link-local address.

RIPng

- Same as IPv4
  - Based on RIPv2
  - Distance vector, max. 15 hop, split-horizon, …
- It's an IPv6 only protocol
  - In a dual-stack environment, running RIP, you'll need RIP (IPv4) and RIPng (IPv6)
- IPv6 related functionality
  - Uses IPv6 for transport
  - IPv6 prefix, next-hop IPv6 address
  - For RIP updates, uses multicast address FF02::9
  - Updates are sent on UDP port 521

RIPng Configuration and Display

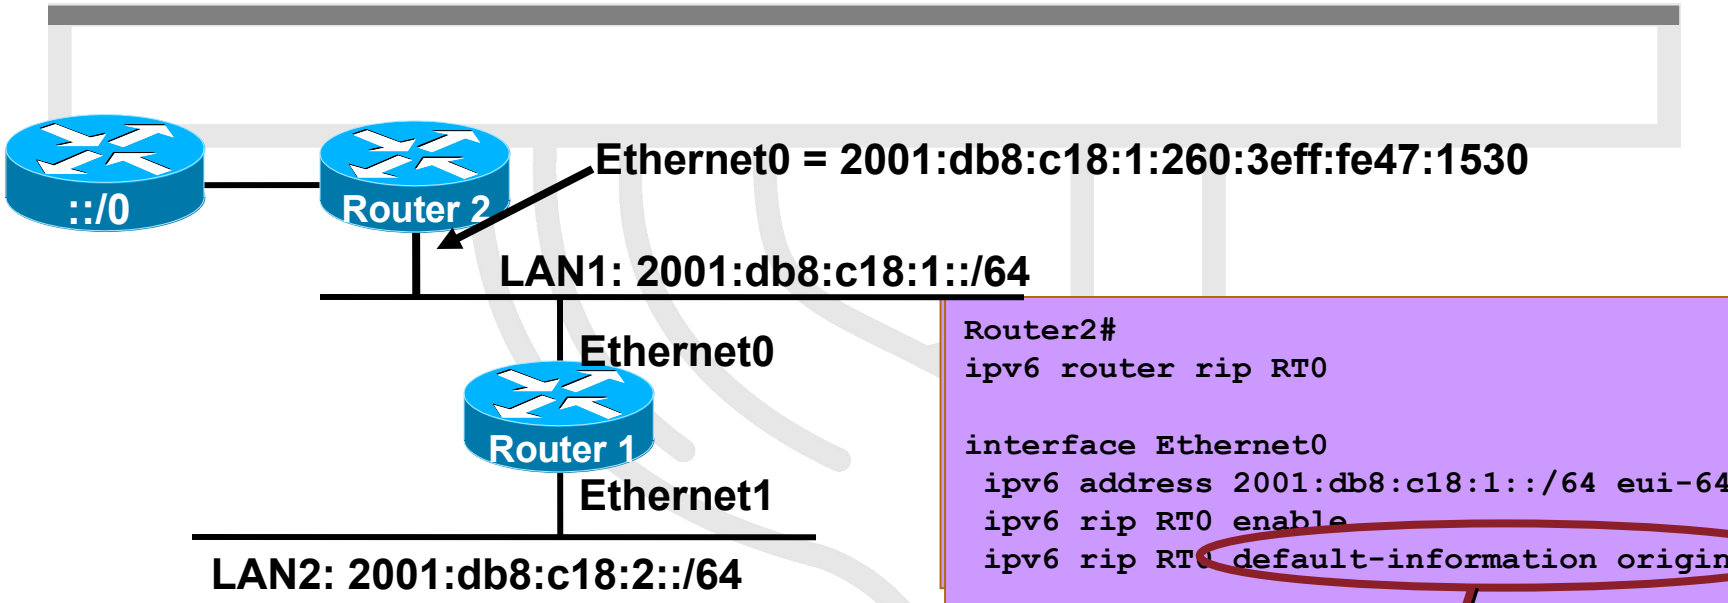


```
Router2#
ipv6 router rip RT0

interface Ethernet0
 ipv6 address 2001:db8:c18:1::/64 eui-64
 ipv6 rip RT0 enable
 ipv6 rip RT0 default-information originate
```

```
Router1#
ipv6 router rip RT0

interface Ethernet0
 ipv6 address 2001:db8:c18:1::/64 eui-64
 ipv6 rip RT0 enable
Interface Ethernet1
 ipv6 address 3ffe:b00:c18:2::/64 eui-64
 ipv6 rip RT0 enable
```

```
Router2# debug ipv6 rip
RIPng: Sending multicast update on Ethernet0 for RT0
       src=FE80::260:3eff:fe47:1530
       dst=FF02::9 (Ethernet0)
       sport=521, dport=521, length=32
       command=2, version=1, mbz=0, #rte=1
       tag=0, metric=1, prefix=::/0
```

Multicast all Rip-Routers

Link-local src address

ISISv6

- OSI Protocol - Originally designed as an intra-domain routing protocol for Connectionless Network Service (CLNS) traffic

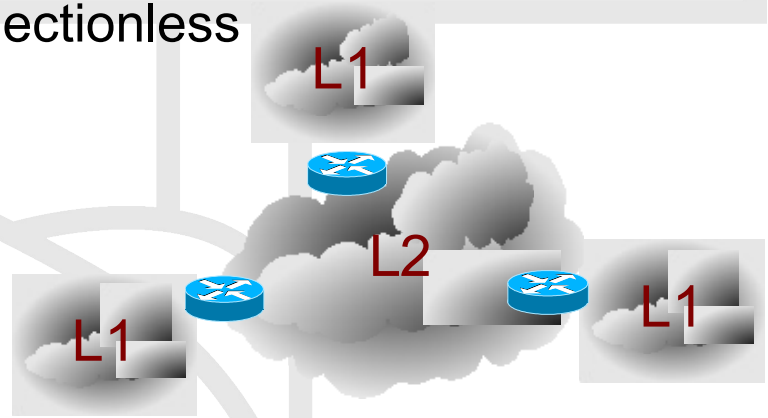

- Based on two levels
  - L2 = Backbone
  - L1 = Stub
  - L2L1= interconnect L2 and L1
- Runs on top of CNLS
  - Each IS device still sends out LSP (Link State Packets)
  - Send information via TLV's (Tag/Length/values)
  - Neighborship process is unchanged
- **Major operation remains unchanged**

ISISv6 #2

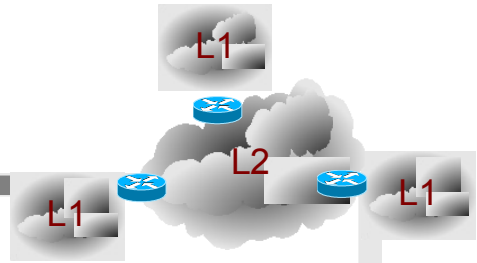

- Updated features:
  - Two new Tag/Length/Values (TLV) for IPv6
    - IPv6 Reachability Describes network reach-ability, contains V6 routing prefix & Metric
    - IPv6 Interface Address
      - » Contains IPv6 interface address (128 bit vs. 32)
      - » For Hello PDUs, must contain the Link-Local address
      - » For LSP, must only contain the non-Link Local address
  - New network Layer Identifier
    - IPv6 NLPID
- Runs on data link. If tunneled, must be mode GRE not IPV6IP

Cisco IOS IS-IS dual IP configuration

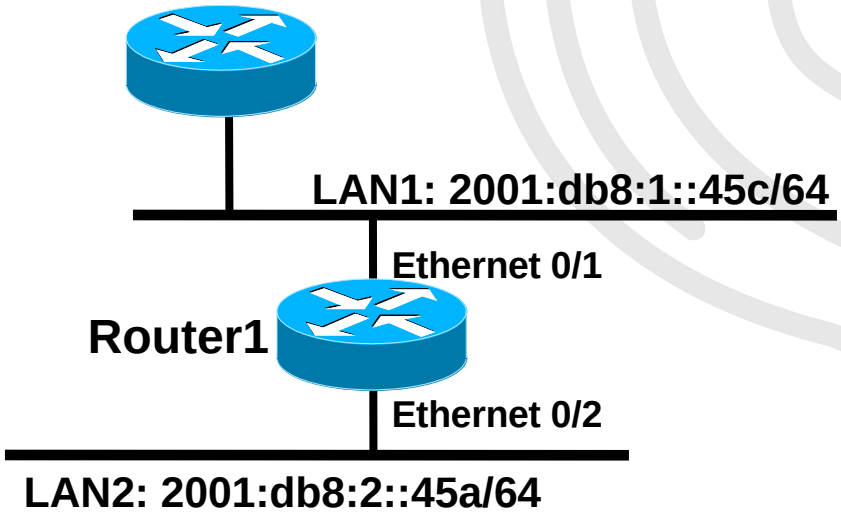

Dual IPv4/IPv6 configuration. Redistributing both IPv6 static routes and IPv4 static routes.

```
Router1#
 interface ethernet0/1
     ip address 10.1.1.1 255.255.255.0
     ipv6 address 2001:db8:1::45c/64
     ip router isis
     ipv6 router isis

  interface ethernet0/2
     ip address 10.2.1.1 255.255.255.0
     ipv6 address 2001:db8:2::45a/64
     ip router isis
     ipv6 router isis

  router isis
      address-family ipv6
      redistribute static
      exit-address-family
      net 42.0001.0000.0000.072c.00
      redistribute static
```

Single SPF rules

- If IS-IS is used for both IPv4 and IPv6 in an area, both protocols must support the same topology within this area.
  - Could set “no adjacency-check” between L2 routers
- All interfaces configured with IS-ISv6 must support IPv6
  - Can’t be configured on MPLS/TE since IS-ISv6 extensions for TE are not yet defined
- All interfaces configured with IS-IS for both protocols must support both of them
  - IPv6 configured tunnel won’t work, GRE should be used in this configuration
- Otherwise, consider Multi-Topology IS-IS (separate SPF)

The problem

Physical Topology **Shortest Path Tree**

- From A perspective, C is only reachable through B

  There is no path from E to C

- All protocols carried by IS-IS have to agree on the same SPT

  No way to distribute traffic across the domain

  All links need to understand all protocols

The need

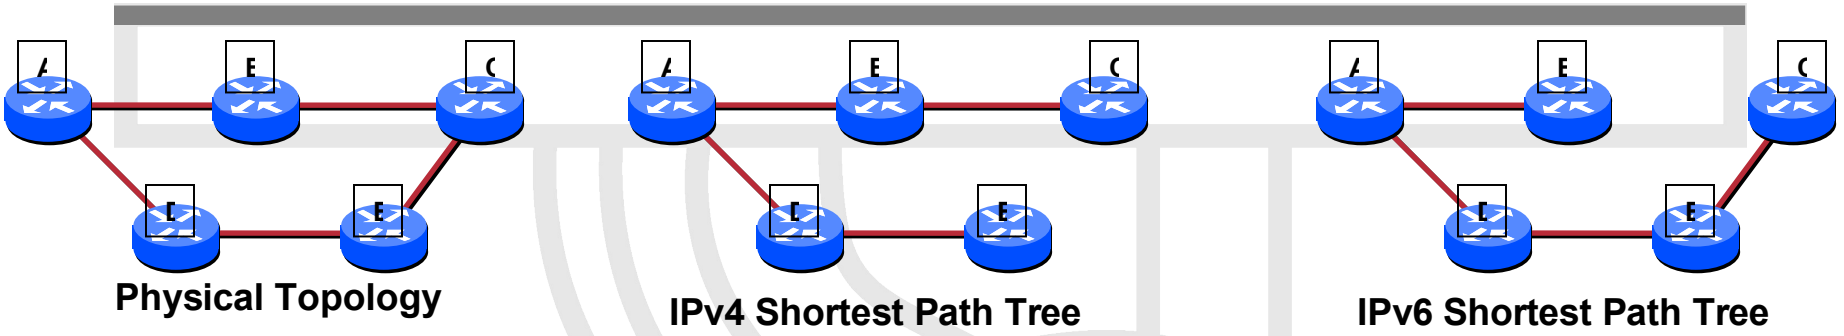

- Ability to distribute traffic across all links

  Separate traffic per address families

Cisco IOS Multi-Topology IS-IS configuration example

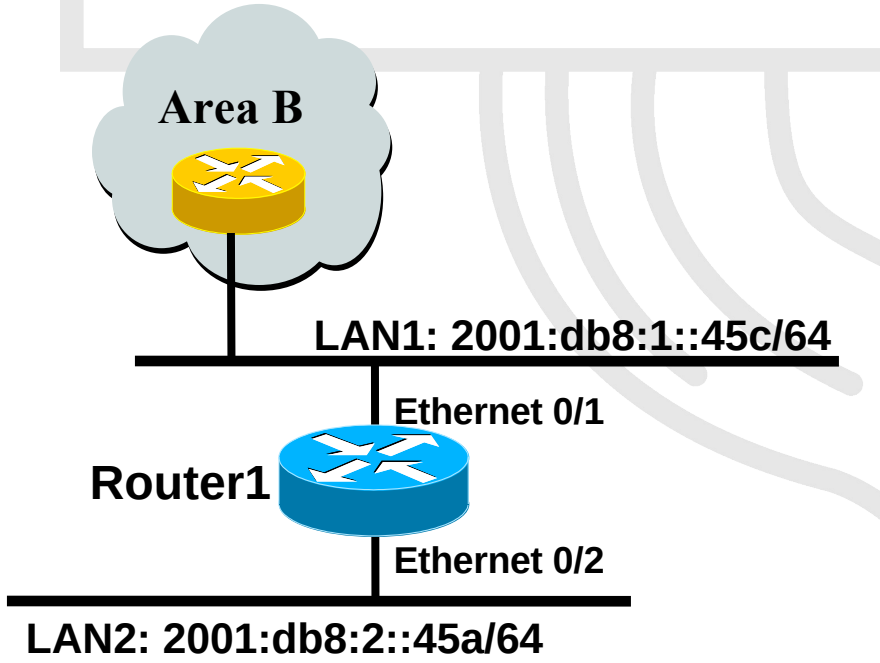


```
Router1#
 interface ethernet 0/1
    ip address 10.1.1.1 255.255.255.0
    ipv6 address 2001:db8:1::45c/64
    ip router isis
    ipv6 router isis
    isis ipv6 metric 20

 interface ethernet 0/2
    ip address 10.2.1.1 255.255.255.0
    ipv6 address 2001:db8:2::45a/64
    ip router isis
    ipv6 router isis
    isis ipv6 metric 20

router isis
net 49.0000.0100.0000.0000.0500
metric-style wide
!
address-family ipv6
multi-topology
exit-address-family
```

- The optional keyword ***transition*** may be used for transitioning existing IS-IS IPv6 single SPF mode to MT IS-IS.
- Wide metric is mandated for Multi-Topology to work.

OSPFv3

- OSPFv3 = OSPF for IPv6 (RFC 2740)
- Based on OSPFv2

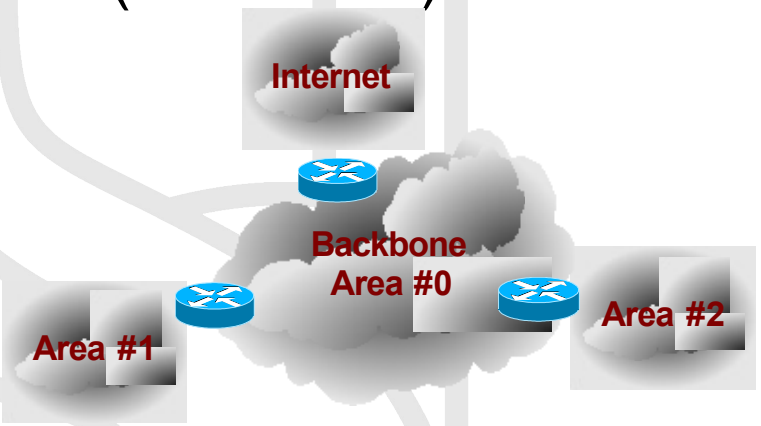

- Topology of an area is invisible from outside the area
  - LSA flooding is bounded by area
  - SPF calculation is performed separately for each area
- All areas must have a connection to the backbone

OSPFv3

- OSPFv3 is an IPv6-only protocol
  - In a dual-stack environment, running OSPF, you'll need OSPFv2 (IPv4) and OSPFv3 (IPv6)
  - There is some work-in-progress about extensible mechanisms to enable OSPFv3 with the support for different address families

- Updated Features
  - Runs directly over IPv6 – you can use link-locals
  - Distributes IPv6 prefixes
  - New LSA types
  - Uses the Multicast address
    - ALLSPFRouters (FF02::5)
    - ALLDRouters (FF02::6)

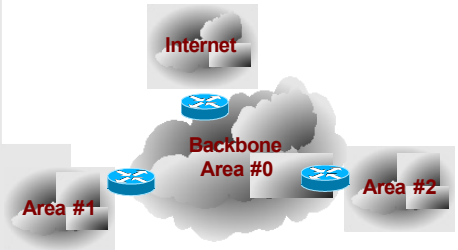

What IPv6 Attributes Affect OSPF?

- 128 bit addresses
- Link-Local address
- Multiple addresses per interface
- Authentication

OSPFv3 / OSPFv2 Differences

- OSPFv3 runs over a link, rather than a subnet
- Multiple instances per link
- OSPFv3 topology not IPv6-specific
  - Router ID
  - Link ID
- Standard authentication mechanisms (IPSec)
- Uses link-local addresses
- Generalized flooding scope

New LSA Types

- Link LSA

  Informs neighbors of link local address

  Informs neighbors of IPv6 prefixes on link

- Intra-Area Prefix LSA

  Associates IPv6 prefixes with a network or router

Removal of Address Semantic

- IPv6 address is not present in OSPFv3 packets

  Exception: LSA payload

- Router-LSA and Network-LSA expressing topology
- Router ID, area ID, LSA link state ID remain a 32 bit number
- Neighbors are always identified by Router ID

ABR Configuration

```
ipv6 unicast-routing
!
interface Ethernet0/0
ipv6 address 2001:db8:1:1::1/64
ipv6 ospf 1 area 0
!
interface Ethernet0/1
ipv6 address 2001:db8:1:2::2/64
ipv6 ospf 1 area 1
!
ipv6 router ospf 1
router-id 2.2.2.2
area 1 range 2001:db8:1:2::/48
```

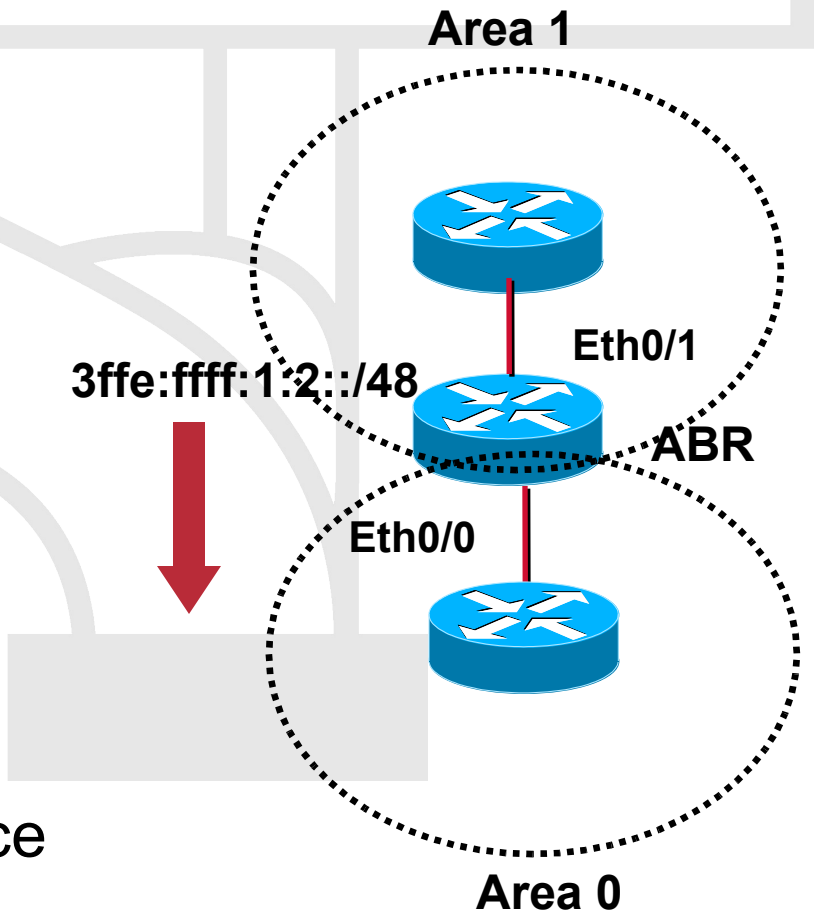

Configuration changed
- IPv4 - specify network
- IPv6 - specification on interface

MP-BGP Basics

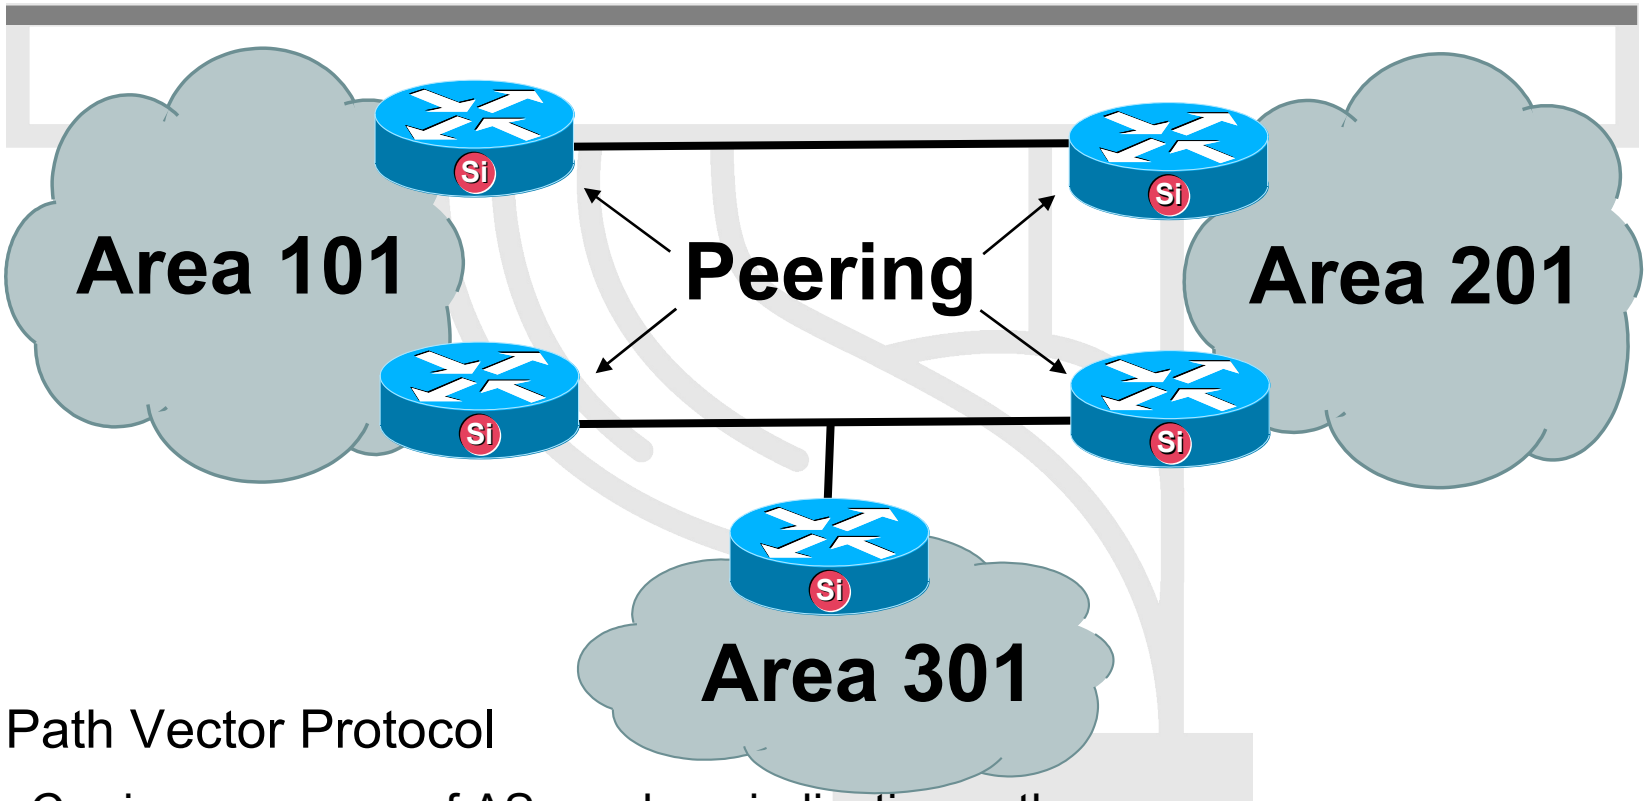

- Path Vector Protocol

  Carries sequence of AS numbers indicating path
- Ties Autonomous Systems together via Peering
- Multiple address families: ipv4, ipv6, unicast, multicast

Multiprocol BGP

- Exterior Gateway Protocol
- Connect separate routing domains that contain independent routing policies (AS)
- Carries sequences of AS numbers indicating path
- Supports the same features and functionality as IPv4 BGP
- Multiple addresses families: IPv4, IPv6, unicast, multicast

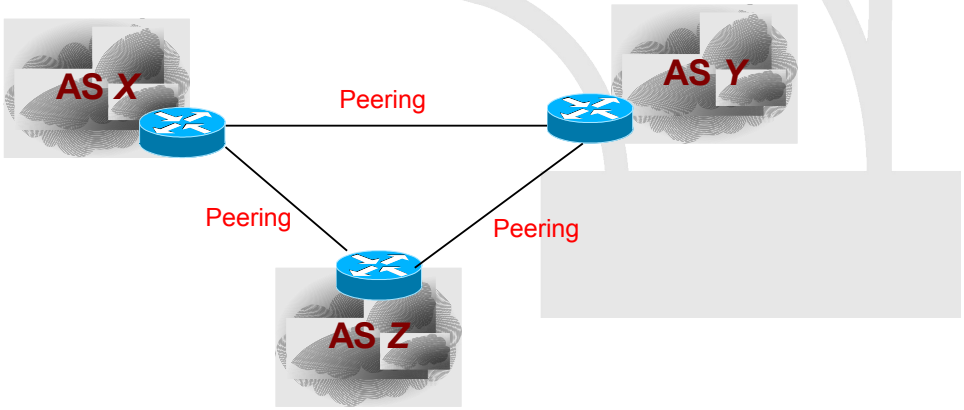

Multiprotocol BGP

- BGP4 carries only 3 types of information wich is truly IPv4 specific:
  - NLRI in the UPDATE message contains an IPv4 prefix
  - NEXT_HOP attribute in the UPDATE message contains an IPv4 address
  - BGP ID in AGGREGATOR attribute

- RFC 2858 defines multi-protocols extensions for BGP4
  - this makes BGP4 available for other network layer protocols (IPv6, MPLS…)
  - New BGP4 attributes (optional, transitive):
    - MP_REACH_NLRI
    - MP_UNREACH_NLRI
    - Messages contains triplets:
      - Address Family Information (AFI)
      - Next-Hop Information (must be of the same address family)
      - NLRI
  - Protocol Independent NEXT_HOP attribute
  - Protocol Independent NLRI attribute

BGP-4 Extensions for IPv6

- Address Family Information (AFI) for IPv6
  - AFI = 2 (RFC 1700)
  - Sub-AFI = 1 Unicast
  - Sub-AFI = 2 (Mulitcast for RPF check)
  - Sub-AFI = 3 for both Unicast and Mulitcast
  - Sub-AFI = 4 Label
  - Sub-AFI= 128 VPN

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

BGP-4 Configurations for IPv6

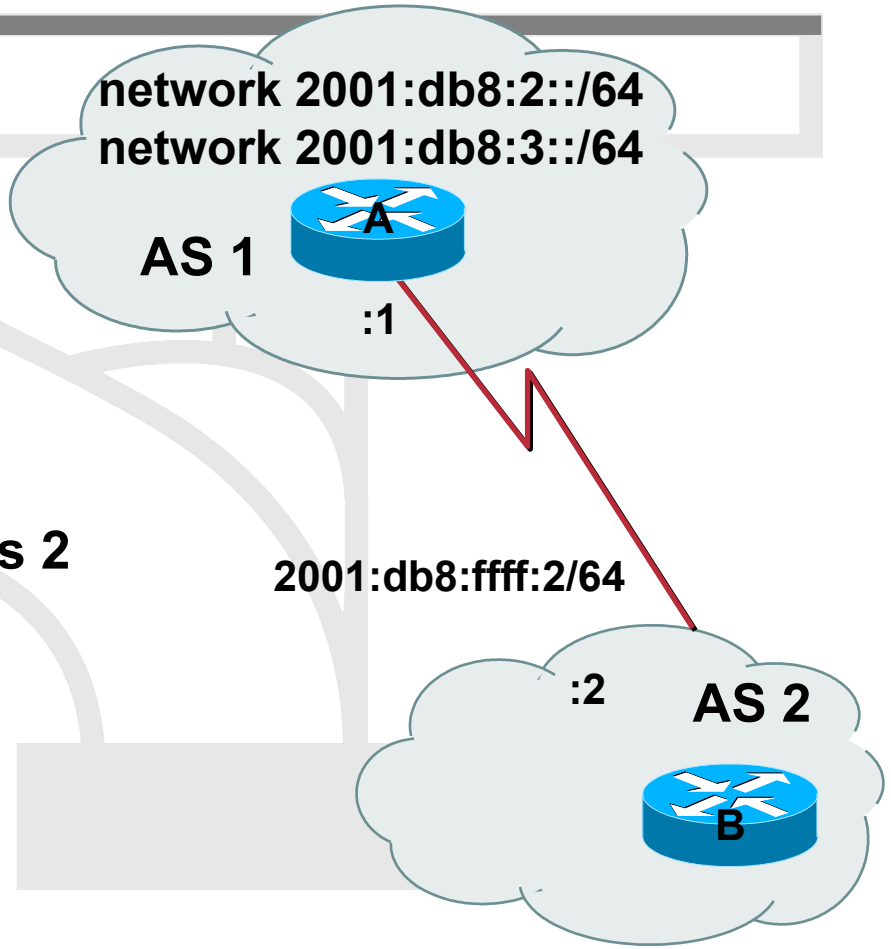

Router A

```
router bgp 1
no bgp default ipv4 unicast
bgp router-id 1.1.1.1
neighbor 2001:db8:ffff:2::2 remote-as 2
address-family ipv6
neighbor 2001:db8:ffff:2::2 activate
 network 2001:db8:2::/64
 network 2001:db8:3::/64
```

Conclusions

- All major routing protocols have stable IPv6 support
- And there isn't major differences with IPv4
- In a dual-stack environment, running OSPF, you'll need OSPFv2 (IPv4) and OSPFv3 (IPv6). It may change in a near future.
- In a dual-stack environment, running RIP, you'll need RIPv1/RIPv2 (IPv4) and RIPng (IPv6)