Routing Protocols

Malta

April 4-6 2006

Jerome.Durand@renater.fr

IPv**6DISS**emination and Exploitation

Agenda

- Internal Routing
  - RIPng
  - IS-IS
  - OSPFv3

- External Routing
 - Multiprotocol BGP

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

IS-IS for IPv6

- OSI Protocol

- Based on two levels
  - L2 = Backbone
  - L1 = Stub
  - L2L1= interconnect L2 and L1

- Runs on top of CNLS (ConnectionLess Network Service)
  - Each IS device still sends out LSP (Link State Packets)
  - Send information via TLV's (Tag/Length/values)
  - Neighborship process is unchanged
  - Requires NSAP addresses to be configured

- **Major operation remains unchanged**

L1

L2

L1

L1
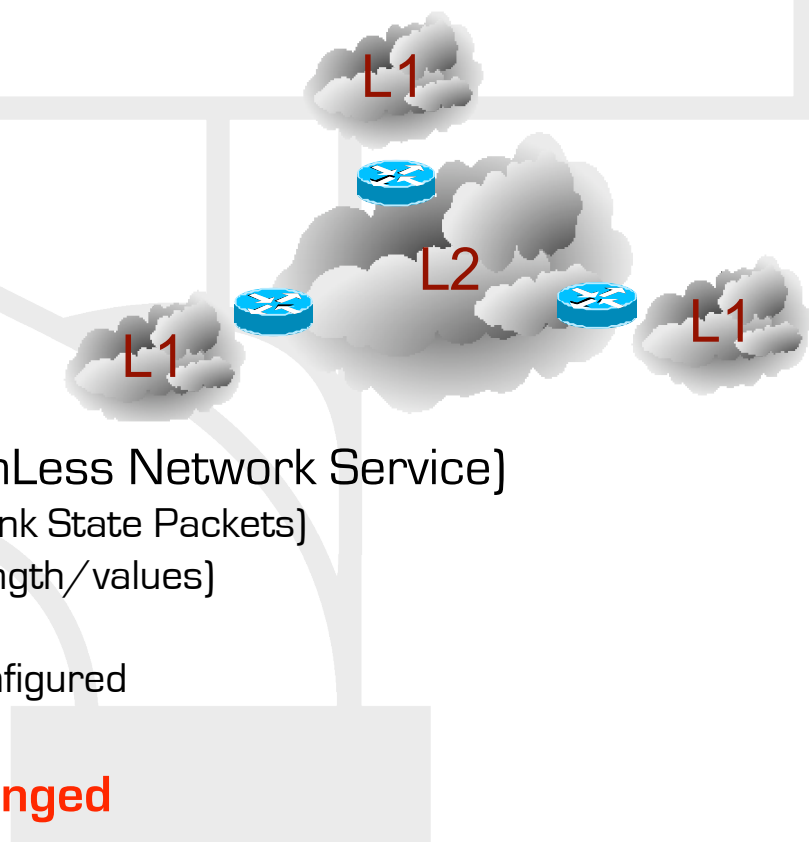
ISIS for IPv6 #2

- Updated features:
  - Two new Tag/Length/Values (TLV) for IPv6
    - IPv6 Reachability
    - IPv6 Interface Address

  - New network Layer Identifier
    - IPv6 NLPID

- A single instance of the protocol is needed in a dual-stack environment.

OSPFv3

- OSPFv3
  - New version of OSPF to support IPv6
- Very similar to OSPFv2
  - Subnets → Links

Internet

**Backbone
Area #0**

Area #1

Area #2

- 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
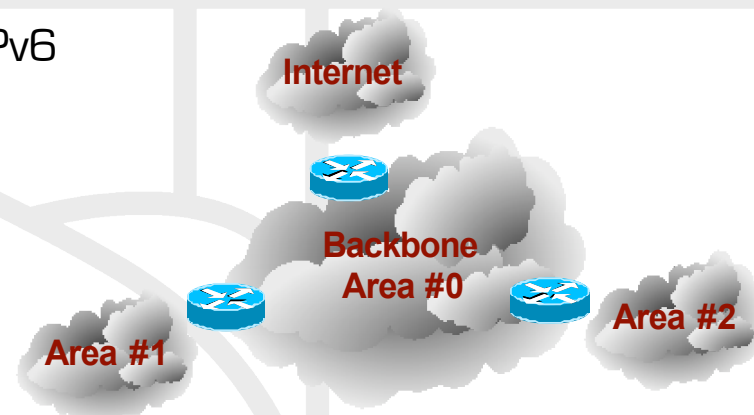
OSPFv3

- OSPFv3 is only implemented for IPv6 at that time
  - 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
  - Transported in IPv6 packets
  - Distributes IPv6 prefixes
  - New LSA types
  - Uses the Multicast address
    - ALLSPFRouters   (FF02::5)
    - ALLDRouters   (FF02::6)
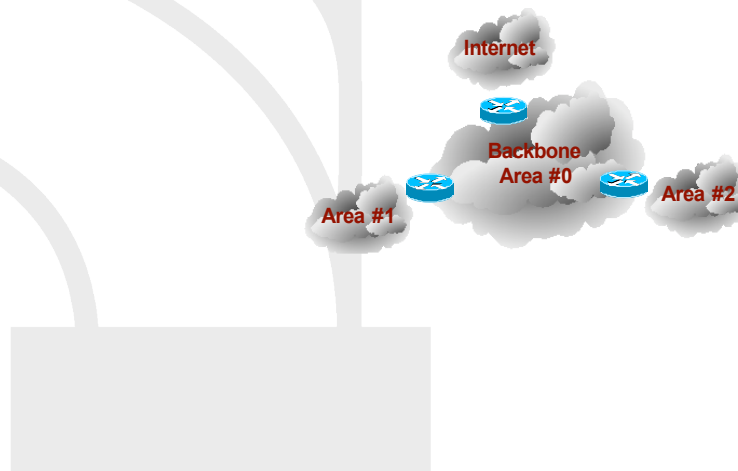
Multiprocol BGP

- Border 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
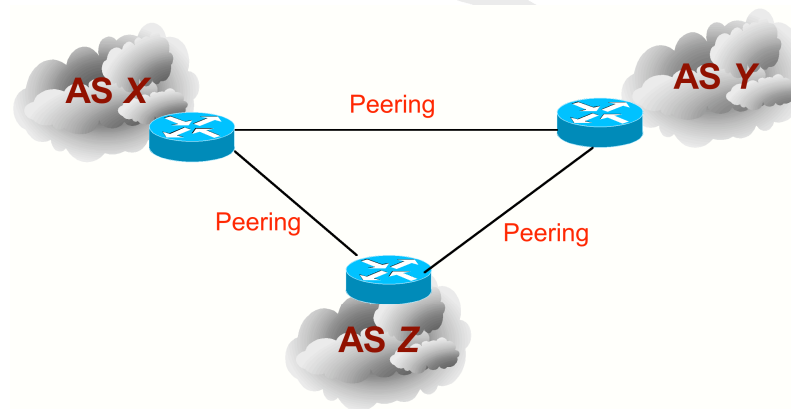
Multiprotocol BGP

- 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:

    - MP_REACH_NLRI

    - MP_UNREACH_NLRI

  - Protocol Independent NEXT_HOP attribute

  - Protocol Independent NLRI attribute

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)