Routing Protocols

Internal and External Routing

6DEPLOY. IPv6 Deployment and Support

Copy Rights

This slide set is the ownership of the 6DEPLOY 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 6DEPLOY courtesy

PDF files are available from www.6deploy.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
- Pedro Lorga, FCCN, Portugal

Contributors

- Mónica Domingues, FCCN, Portugal
- Paulo Ferreira, FCCN, Portugal

Prerequisites

You should have followed previously the modules:

- 010-IPv6 Introduction
- 020-IPv6 Protocol
- 030-IPv6 Addressing

Agenda

Gateway Redundancy

- VRRP

Internal Routing

- RIPng
- IS-IS
- OSPFv3

External Routing

- Multiprotocol BGP

VRRP

IETF: Version 3

- RFC5798, March 2010
- Based on VRRPv2 for IPv4
- Election protocol

Usage of «virtual» addresses

- Which are used by/configured on hosts
- One of the existent VRRP routers is elected as «MASTER»

IPv6 Multicast Address

- Assigned by IANA = FF02::12

VRRP

Advantage of using VRRP on IPv4:

- Higher-availability default path without requiring configuration of dynamic routing or router discovery protocols on every end-host.

Advantage of using VRRP on IPv6:

- Quicker switchover to Backup routers than can be obtained with standard IPv6 Neighbor Discovery mechanisms.

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

ISISv6

OSI Protocol

Based on two levels

- L2 = Backbone
- L1 = Stub
- L2L1= interconnect L2 and L1

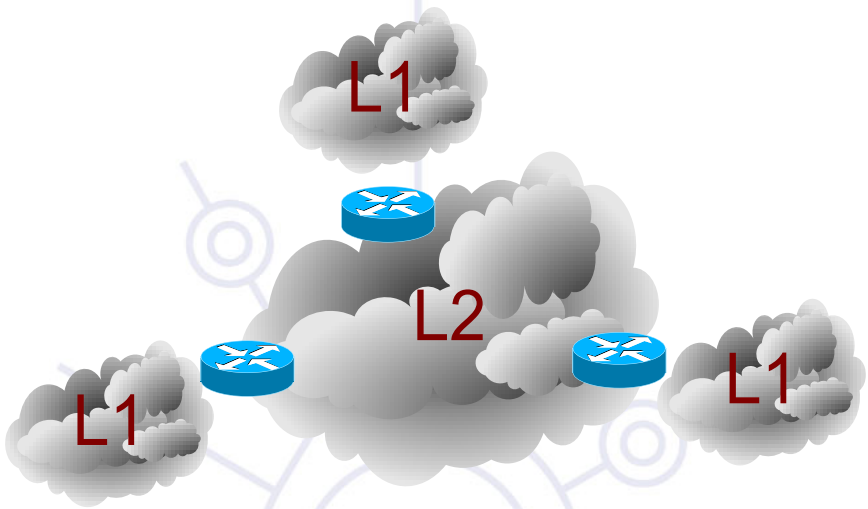

Runs on top of CNLS

- 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
  - IPv6 Interface Address
- New network Layer Identifier
  - IPv6 NLPID

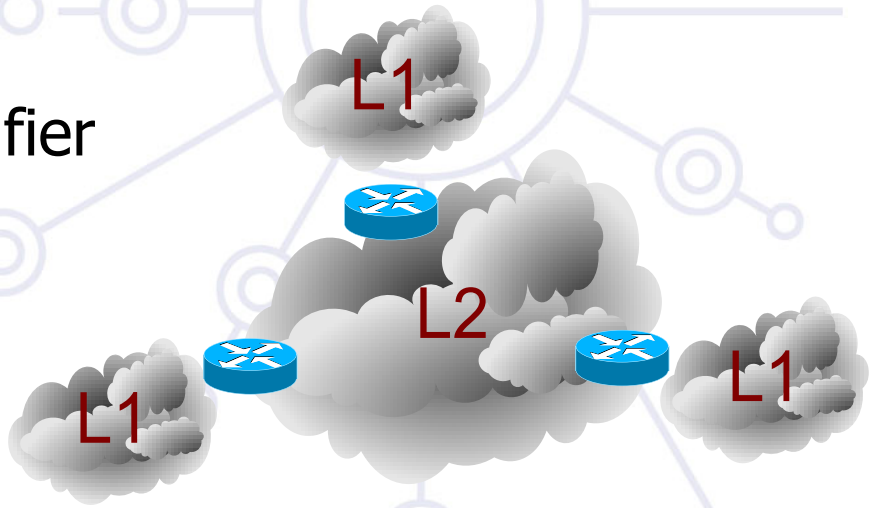

OSPFv3

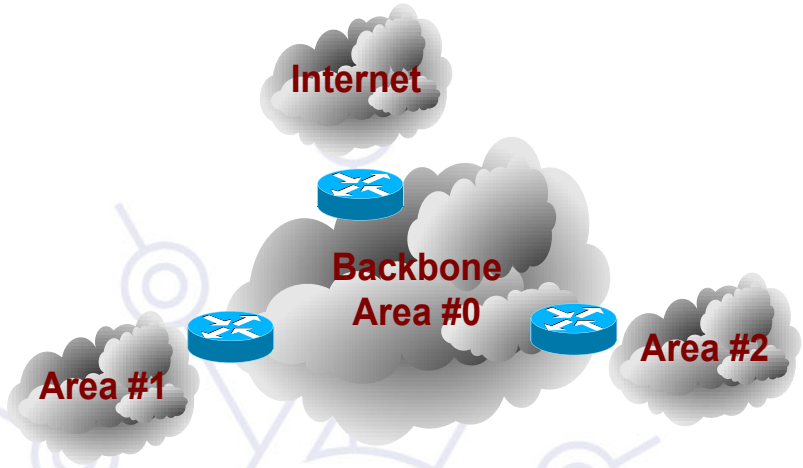

OSPFv3 = OSPF for IPv6

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)
- Work-in-progress about extensible mechanisms to enable OSPFv3 with different address families support

Updated Features

- Runs directly over IPv6
- Distributes IPv6 prefixes
- New LSA types
- Uses Multicast addresses
  - ALLSPFRouters (FF02::5)
  - ALLDRouters (FF02::6)

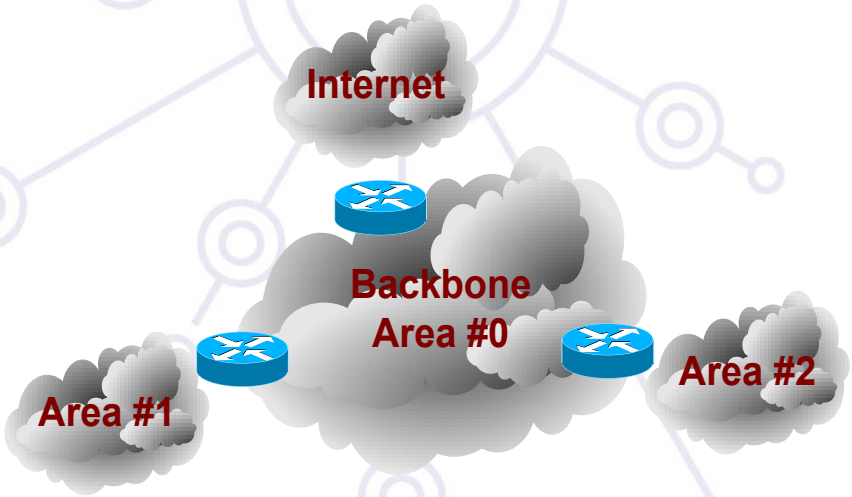

Multiprotocol BGP

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

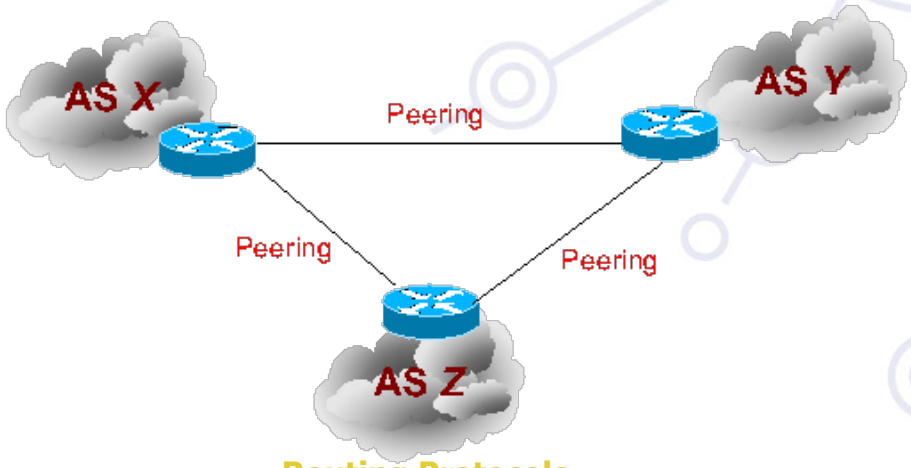

Multiprotocol BGP

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

Multiprotocol BGP

RFC 4760 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 no 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)

Questions?

Extra Slides

Routing (on systems)

There is always an IPv4 and an IPv6 routing context in every system.

OS	IPv4	IPv6
Cisco (IOS)	**show ip route**	**show ipv6 route**
WinXP	**route print**	**netsh interface ipv6 show route**
Linux	**/sbin/route**	**/sbin/route –A inet6**

Routing Stats (IPv6 vs. IPv4, globally)

(11/09/2008)	IPv6	IPv4
ROUTES	**1505**	**281136**
AGGREGATED ROUTES	**1400 (93,02%)**	**170595 (60,68%)**
AUTONOMOUS SYSTEMS	**1131**	**29345**

source: www.cidr-report.org