IPv6 Tutorial

János Mohácsi

Networks 2008 - IPv6 tutorial Budapest, 28 Sept 2008

6DEPLOY Partners

NRENs

Renater France
GRNET Greece
FCCN Portugal
NIIF/HUNGARNET Hungary
UNINETT Norway
BREN Bulgaria

SMEs

Consulintel Spain
Martel Consulting (coordinator) Switzerland

Industry

Cisco Netherlands

Universities

UCL United Kingdom
Soton ECS United Kingdom

Non-European Partners

AfriNIC Mauritius
LACNIC Uruguay

Associated partners: RIPE NCC, APNIC

6deploy.org

Project Objectives

Support of EU policy

The Internet is now the main telecommunications technology that underpins all aspects of business and leisure, and as such is central to the economic growth of a country. Awareness of the evolution of the Internet, and providing support for the introduction of IPv6 is therefore crucial as ICT becomes a major theme in FP7

Specific ***technical focus*** on supporting the deployment of IPv6 in:

- research infrastructures, for supporting all fields of science and technology
- FP7 projects (especially in the areas of emergency services, healthcare, transport, gaming)
- developing countries (Africa, Latin America, Asia and E. Europe), and
- commercial organisations in Europe

Networks 2008 - IPv6 tutorial

6deploy.org

History

FP4: 6INIT, 6WINIT

FP5: 6NET, Euro6IX, Occasion, 6Power, IPv6 TF

FP6: 6DISS, Sponge, IPv6 TF (continued)

FP7: 6DEPLOY

6DEPLOY is the one we exploit the most, in terms of partners and material

Networks 2008 - IPv6 tutorial

Technical Approach

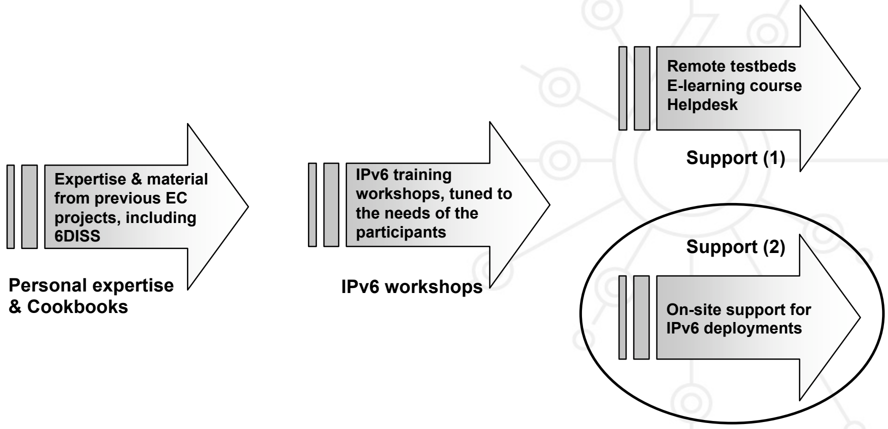

Networks 2008 - IPv6 tutorial

The 6DEPLOY Toolkit

- Workshops for direct training, and for „training other trainers“
- Presentation material on 20 topics associated with IPv6
- Practical configuration exercises
- Professional e-learning package
- Remote testbeds in Paris and Sofia (for use in- and out- side the workshops)
- Book on deployment guidelines (from 6NET)
- Helpdesk service run by experienced persons
- Website with links to 6DEPLOY documents and external sources

www.6diss.org/e-learning
www.6deploy.org

Copy …Rights

Contributions

Main authors

- János Mohácsi, NIIF/HUNGARNET - Hungary
- Jérôme Durand, Bernard Tuy, Renater, France
- Tim Chown, University of Southampton, Great-Britain
- Stig Venaas, UNINETT
- Joao Nuno Ferreira, Carlos Friacas, FCCN

Why a new version for IP ?

Historical facts

- 1983 : Research network for ~ 100 computers
- 1992 : Commercial activity
- Exponential growth
- 1993 : Exhaustion of the class B address space
- Forecast of network collapse for 1994!

6deploy.org

IPv4 address space consumption

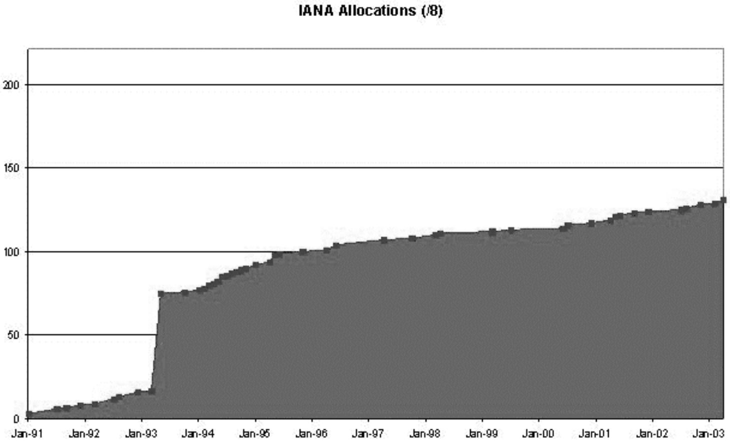

Networks 2008 - IPv6 tutorial

6deploy.org

IPv4 address space consumption /2

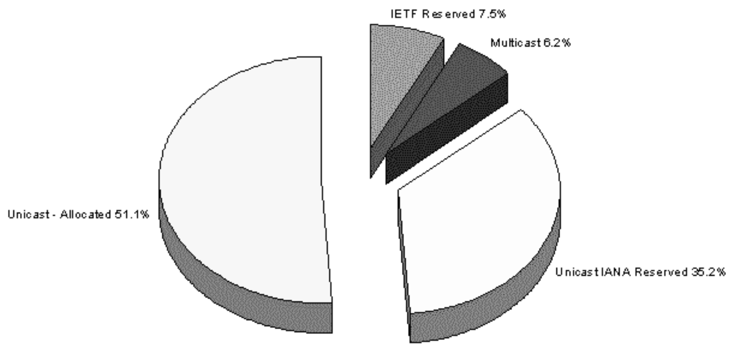

Networks 2008 - IPv6 tutorial

6deploy.org

IPv4 prefixes consumption pace

Year	Month	available /8s (IANA)	Yearly consumption
2006	September	59	
	December	55	12
2007	September	44	15
	December	42	13
2008	June	39 (-2) ?	6 (10) ?

6deploy.org

IPv4 address space depletion

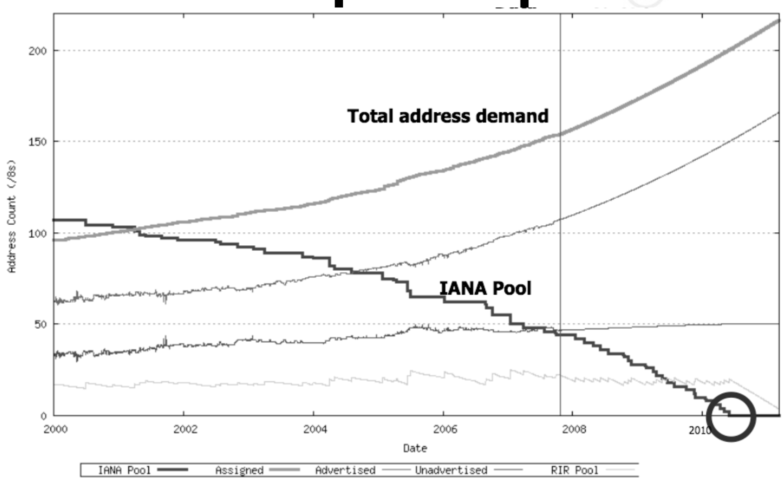

Geoff Huston
APNIC
Sept. 2007

6deploy.org

Emergency Measures

CIDR
Private addresses
NAT

Montpellier, Fr -Octobre 2008 15
Networks 2008 - IPv6 tutorial

6deploy.org

CIDR …

Allocate exceptionally class B addresses
Re-use class C address space
CIDR (*Classless Internet Domain Routing*)

- RFC 1519 (PS)
- network address = prefix/prefix length
- Classes abandon = less address waste
- allows aggregation (reduces routing table size)

Montpellier, Fr -Octobre 2008 16
Networks 2008 - IPv6 tutorial

Private addresses
(RFC 1918)

Allow private addressing plans
Addresses are used internally
Similar to security architecture with firewall
Use of proxies or NAT to go outside
- RFC 1631, 2663 and 2993

NAT-PT
- the most commonly used of NAT variations in the IPv6 world

NAT (continued)

Advantages:
- Reduce the need of official addresses
- Ease the internal addressing plan
- Transparent to some applications
- "Security" vs obscurity
- Netadmins/sysadmin

Disadvantages:
- Translation sometime complex (e.g. FTP)
- Apps using dynamic ports
- Does not scale
- Introduce states inside the network:
  - Multihomed networks
- Breaks the end-to-end paradigm
- Security with IPsec

=> Should be reserved for small sites in Client/Server mode

Emergency Measures

These emergency measures gave time to develop a new version of IP, named IPv6

IPv6 keeps principles that have made the success of IP

Corrects what was wrong with the current version (v4)

BUT are emergency measures enough?

From emergency to IPv6

IPv6 is already there …

- Internet v6 is there today :
- NRENs in EU, North America, Asia … are interconnected in IPv6
- Lots of IXP are offering IPv6 connectivity
- ISPs and Telcos exchange IPv6 routes

Then the question is not "if" but "when ?" and "how ?"

By Sep. 11th 2008 resources exhaustion are projected

- IANA pool : Oct. 2010
- RIRs pool : Oct. 2011
- Data from : http://www.potaroo.net/tools/ipv4/index.html

IPv6 Protocol (RFC 2460 DS)

- IPv6 Header
- IPv6 addressing
- Protocols associated to IPv6

Networks 2008 - IPv6 tutorial Budapest, 28 Sept 2008

6deploy.org

IPv6 Header

Networks 2008 - IPv6 tutorial

IPv4 Header

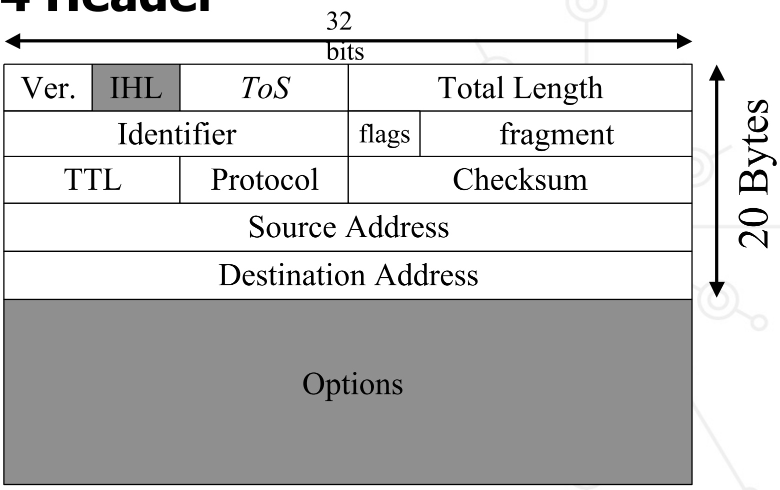

IPv4 Header

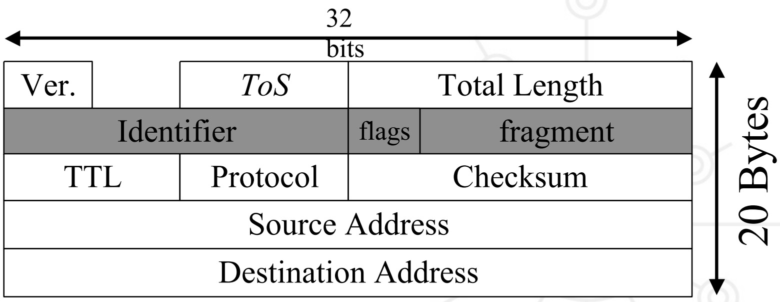

IPv4 Header

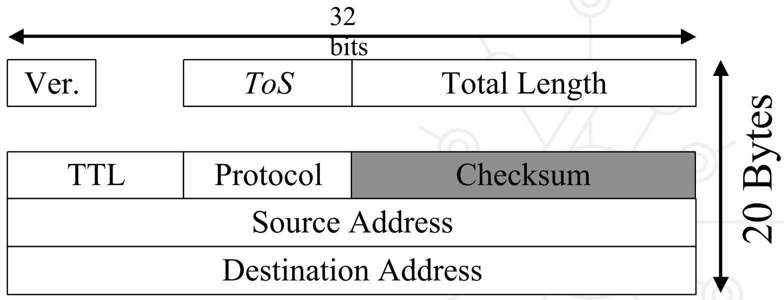

Networks 2008 - IPv6 tutorial

IPv6: Header simplification

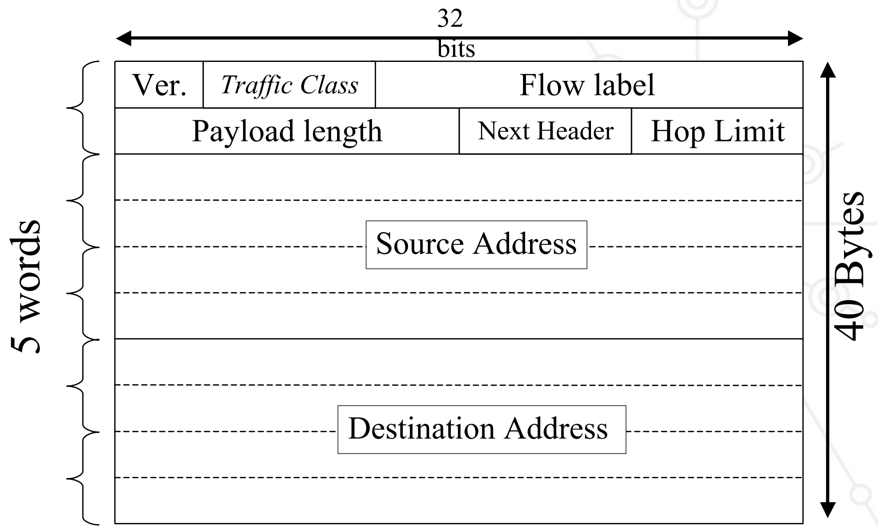

Networks 2008 - IPv6 tutorial

Is it enough for the future ?

Address length

- Between 1 564 and 3 911 873 538 269 506 102 addresses by m^2
- Justification of a fix address length

Hop Limit

- Should not be a problem

Payload Length

- Use Jumbogram for specific cases

6deploy.org

IPv6 extensions

Networks 2008 - IPv6 tutorial

IPv6: Optional headers

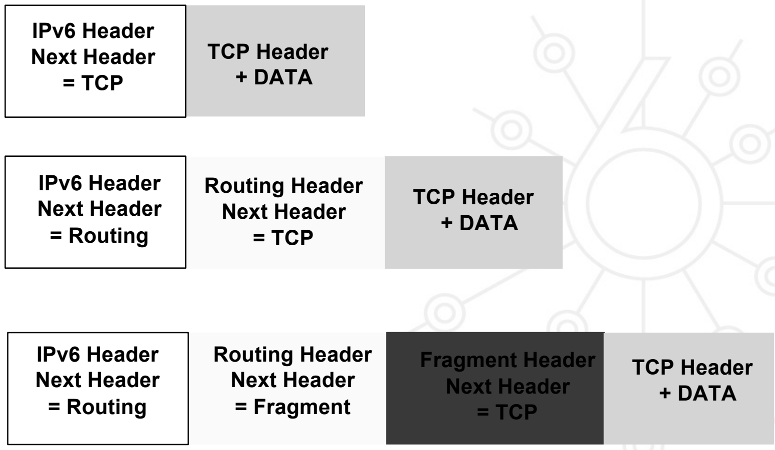

IPv6: Optional extensions

Hop-by-hop (jumbogram, router alert)
- Always the first extension
- Replace IPv4 options,
- Analyzed by every router.

Destination

Routing (loose source routing)

Fragmentation

Authentication

Security

6deploy.org

v4 options vs. v6 extensions

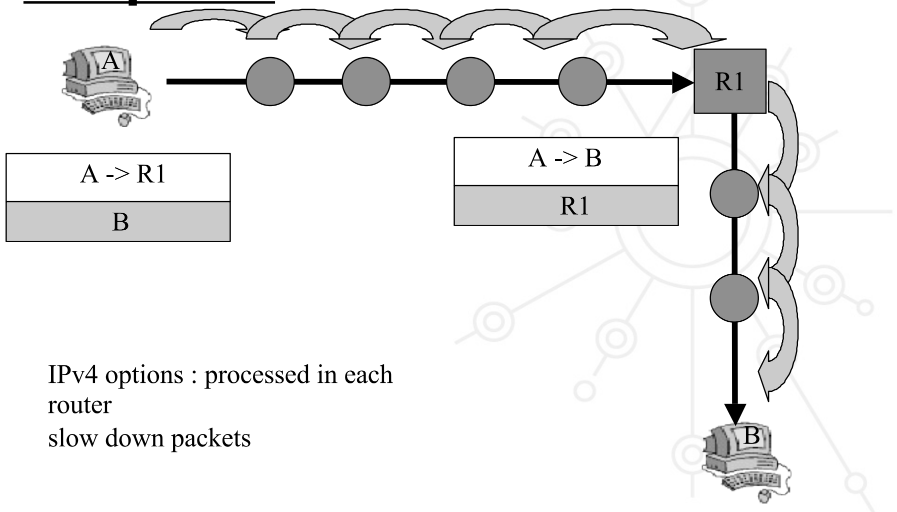

IPv4 options : processed in each router
slow down packets

6deploy.org

v4 options vs. v6 extensions

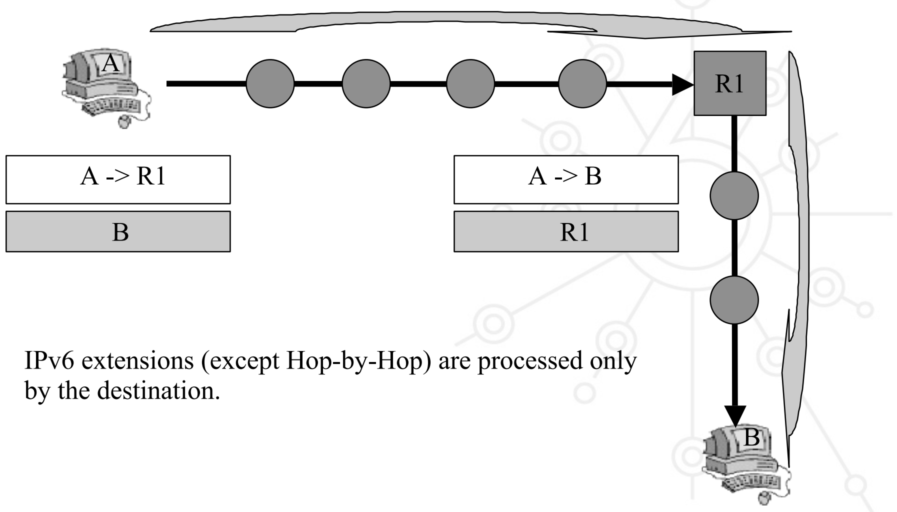

IPv6 extensions (except Hop-by-Hop) are processed only by the destination.

6deploy.org

Order is important

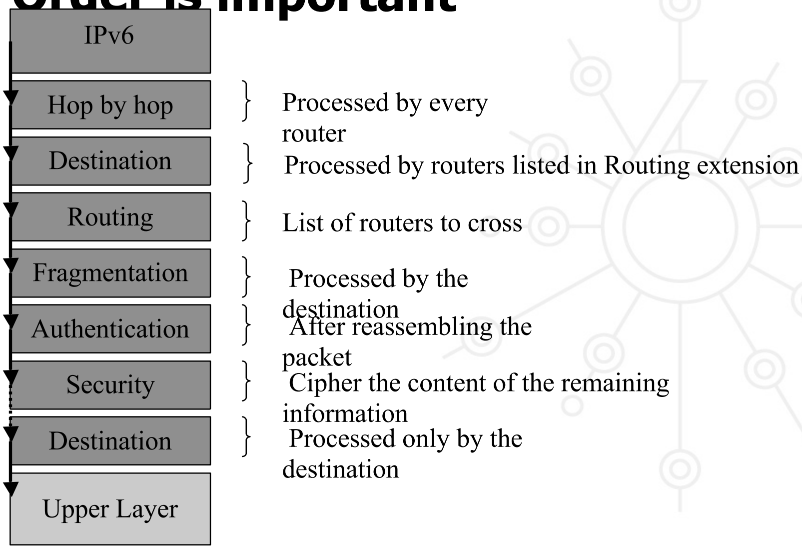

Networks 2008 - IPv6 tutorial

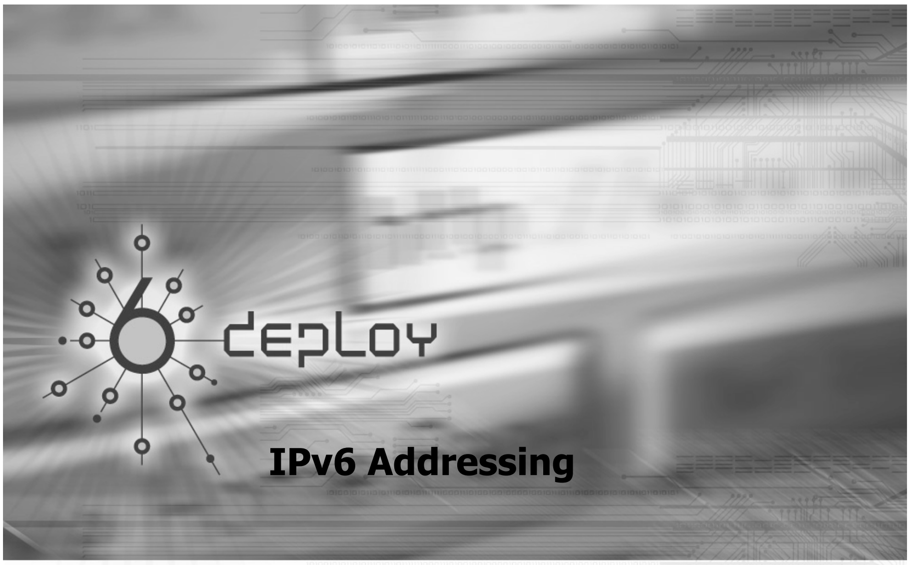

IPv6 Addressing

6DEPLOY. IPv6 Deployment and Support

Networks 2008 - IPv6 tutorial Budapest, 28 Sept 2008

IPv6 Addressing Scheme

RFC4291 defines IPv6 addressing scheme
RFC3587 defines IPv6 global unicast address format
128 bit long addresses
- Allow hierarchy
- Flexibility for network evolutions

Use CIDR principles:
- Prefix / prefix length
  - 2001:660:3003::/48
  - 2001:660:3003:2:a00:20ff:fe18:964c/64
- Aggregation reduces routing table size

Hexadecimal representation
Interfaces have several IPv6 addresses

IPv6 Address Types

Unicast (one-to-one)
- global
- link-local
- site-local (deprecated)
- Unique Local (ULA)
- IPv4-compatible (deprecated)
- IPv6-mapped

Multicast (one-to-many)
Anycast (one-to-nearest)
Reserved

Textual Address Format

Preferred Form (a 16-byte Global IPv6 Address):

2001:0DB8:3003:0001:0000:0000:6543:210F

Compact Format:

2001:DB8:3003:1::6543:210F

IPv4-mapped: ::FFFF:13.1.68.3

Literal representation

- [2001:DB8:3003:2:a00:20ff:fe18:964c]
- http://[2001:DB8::43]:80/index.html

IPv6 Address Type Prefixes

Address Type	Binary Prefix	IPv6 Notation
Unspecified	00…0 (128 bits)	::/128
Loopback	00…1 (128 bits)	::1/128
Multicast	1111 1111	FF00::/8
Link-Local Unicast	1111 1110 10	FE80::/10
ULA	1111 110	FC00::/7
Global Unicast	(everything else)	
IPv4-mapped	00…0:1111 1111:IPv4	::FFFF:IPv4/128
Site-Local Unicast (deprecated)	1111 1110 11	FEC0::/10
IPv4-compatible (deprecated)	00…0 (96 bits)	::IPv4/128

Global Unicast assigments actually use 2000::/3 (001 prefix)
Anycast addresses allocated from unicast prefixes

6deploy.org

IPv6 Address Space

Aggregatable Global Unicast Addresses (001): 1/8
Unique Local Unicast addresses (1111 1110 00): 1/128
Link-Local Unicast Addresses (1111 1110 10): 1/1024
Multicast Addresses (1111 1111): 1/256

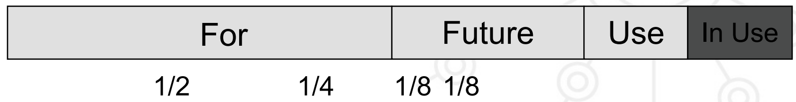

More info:
http://www.iana.org/assignments/ipv6-address-space

Networks 2008 - IPv6 tutorial

6deploy.org

Some Special-Purpose Unicast Addresses

Listed in RFC5156

The **unspecified address,** used as a placeholder when no address is available:

0:0:0:0:0:0:0:0 (::/128)

The **loopback address,** for sending packets to itself:

0:0:0:0:0:0:0:1 (::1/128)

The **documentation prefix [RFC3849]:**
2001:db8::/32

Networks 2008 - IPv6 tutorial

6deploy.org

Link-Local & Site-Local Unicast Addresses

Link-local addresses for use during auto-configuration and when no routers are present **(FE80::/10):**

10 bits	*54 bits*	*64 bits*
1111111010	0 0	Interface ID

FE80

Site-local addresses for independence from changes of TLA / NLA* **(FEC0::/10):** (deprecated by RFC3879)

10 bits	*54 bits*	*64 bits*
1111111011	Subnet ID	Interface ID

FEC0

Networks 2008 - IPv6 tutorial

6deploy.org

Interface IDs

The lowest-order 64-bit field of unicast addresses may be assigned in several different ways:

- auto-configured from a 64-bit MAC address
- auto-configured from a 48-bit MAC address (e.g., Ethernet) expanded into a 64-bit EUI-64 format
- assigned via DHCP
- manually configured
- auto-generated pseudo-random number (to counter some privacy concerns)
- possibly other methods in the future

Network ID	Interface ID
64 bits	64 bits

Networks 2008 - IPv6 tutorial

Autoconfigured Inteface IDs (1)

- **64 bits to be compatible with IEEE 1394 (FireWire)**
- **Eases auto-configuration**
- **IEEE defines the mechanism to create an EUI-64 from IEEE 802 MAC addresses (Ethernet, FDDI)**

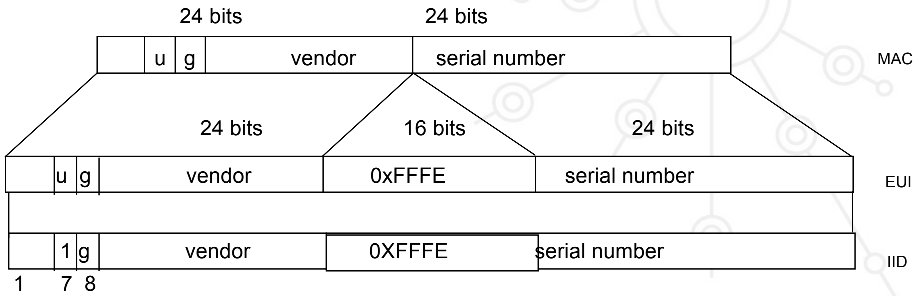

Autoconfigured Inteface IDs (2)

- **Links with non global identifier (e.g., the Localtalk 8 bit node identifier) → fill first left bits with 0**
- **For links without identifiers, there are different ways to proceed (e.g., tunnels, PPP) to have a subnet-prefix-unique identifier:**
  - Choose the universal identifier of another interface
  - Manual configuration
  - Node Serial Number
  - Other Node-Specific Token

Multicast Addresses

11111111	flags	scope	group ID
8	4	4	112 bits

Flags: 0RPT: The high-order flag is reserved, and must be initialized to 0.

- **T**: Transient, or not, assignment
- **P**: Assigned, or not, based on network prefix
- **R**: Rendezvous Point Address embedded, or not

Scope field:
1 - Interface-Local
2 - link-local
4 - admin-local
5 - site-local
8 - organization-local
E - global

(3,F reserved)(6,7,9,A,B,C,D unassigned)

Unique Local IPv6 Unicast Addresses (1)

Defined in **RFC4193:**

- Globally unique prefix with high probability of uniqueness
- Intended for local communications, usually inside a site
- They are not expected to be routable on the Global Internet
- They are routable inside of a more limited area such as a site
- They may also be routed between a limited set of sites
- Locally-Assigned Local addresses vs. Centrally-Assigned Local addresses

6deploy.org

Unique Local IPv6 Unicast Addresses (2)

ULA characteristics:

- Well-known prefix to allow for easy filtering at site boundaries
- ISP independent and can be used for communications inside of a site without having any permanent or intermittent Internet connectivity
- If accidentally leaked outside of a site via routing or DNS, there is no conflict with any other addresses
- In practice, applications may treat these addresses like global scoped addresses

Networks 2008 - IPv6 tutorial

6deploy.org

Unique Local IPv6 Unicast Addresses (3)

Format:

Prefix	L	global ID	subnet ID	interface ID
7 bits	1	40 bits	16 bits	64 bits

FC00::/7 Prefix identifies the Local IPv6 unicast addresses

L = 1 if the prefix is **locally assigned**

L = 0 may be defined in the future (in practice used for **centrally assigned** prefixes)

ULA are created using a pseudo-randomly allocated global ID

- This ensures that there is not any relationship between allocations and clarifies that these prefixes are not intended to be routed globally

Networks 2008 - IPv6 tutorial

6deploy.org

Global Unicast Addresses

Defined in RFC3587

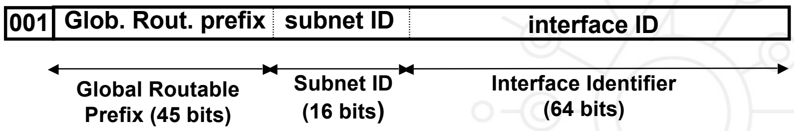

The global routing prefix is a value assigned to a zone (site, a set of subnetworks/links)

- It has been designed as an hierarchical structure from the Global Routing perspective

The subnetwork ID, identifies a subnetwork within a site

- Has been designed to be an hierarchical structure from the site administrator perspective

6deploy.org

Anycast Addresses

Identifier for a set of interfaces (typically in different nodes). A packet sent to an anycast address is delivered to one of the interfaces identified by that address (the "nearest" one, according to the routing protocols' distance)

Taken from the unicast address spaces (of any scope) and are **not syntactically distinguishable from unicast addresses**

Reserved anycast addresses are defined in **RFC 2526**

6Bone Global Unicast Addresses

Obsoleted 6/6/6 – RFC3701

001	TLA	pTLA	NLA*	SLA*	interface ID
	13	12	20	16	64 bits

6Bone: experimental IPv6 network used for testing only

TLA 1FFE (hex) assigned to the 6Bone

- 6Bone addresses start with 3FFE (binary 001 + 1 1111 1111 1110)

Not to be used for production IPv6 service

6to4 Addresses

Defined in **RFC3056: Connection of IPv6 Domains via IPv4 Clouds**

Assigned Prefix: 2002::/16

To assign to sites 2002:V4ADDR::/48

001	0x0002	V4ADDR	SLA	interface ID
6to4 Prefix 2002::/16 (16 bits)		Site Public IPv4 (32 bits)	Site Topology (16 bits)	(64 bits)

Production Addressing Scheme (1)

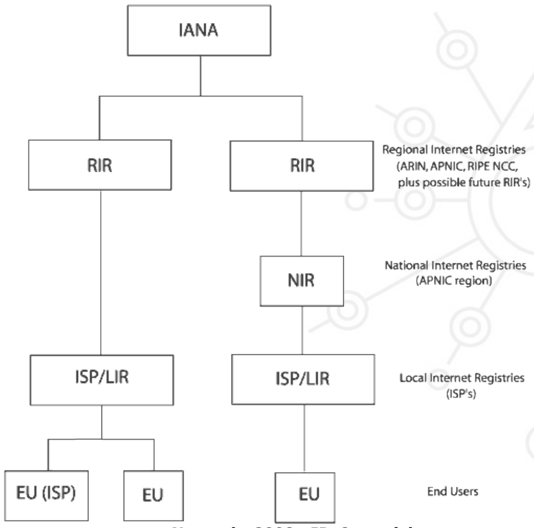

Production Addressing Scheme (2)

001	Glob. Rout. prefix	subnet ID	interface ID
	Global Routable Prefix (45 bits)	Subnet ID (16 bits)	Interface Identifier (64 bits)

LIRs receive by default /32

- Production addresses today are from prefixes 2001, 2003, 2400, etc.
- Can request for more if justified

/48 used only within the LIR network, with some exceptions for critical infrastructures

/48 to /128 is delegated to end users

- Recommendations following RFC3177 and current policies
- /48 general case, /47 if justified for bigger networks
- /64 if one and only one network is required
- /128 if it is sure that one and only one device is going to be connected

Production Addressing Scheme (3)

Source:
http://www.iana.org/assignments/ipv6-unicast-address-assignments

```
IPv6 Global Unicast Address Assignments [0]
[last updated 2008-05-13]
Global Unicast Prefix Assignment     Date       Note
--------------------- ----------     ------     ----
2001:0000::/23        IANA           01 Jul 99  [1]
2001:0200::/23        APNIC          01 Jul 99
2001:0400::/23        ARIN           01 Jul 99
2001:0600::/23        RIPE NCC       01 Jul 99
2001:0800::/23        RIPE NCC       01 May 02
2001:0A00::/23        RIPE NCC       02 Nov 02
2001:0C00::/23        APNIC          01 May 02  [2]
2001:0E00::/23        APNIC          01 Jan 03
2001:1200::/23        LACNIC         01 Nov 02
```

Networks 2008 - IPv6 tutorial

Production Addressing Scheme (4)

```
...
2001:4C00::/23        RIPE NCC       17 Dec 04
2001:5000::/20        RIPE NCC       10 Sep 04
2001:8000::/19        APNIC          30 Nov 04
2001:A000::/20        APNIC          30 Nov 04
2001:B000::/20        APNIC          08 Mar 06
2002:0000::/16        6to4           01 Feb 01  [4]
2003:0000::/18        RIPE NCC       12 Jan 05
2400:0000::/12        APNIC          03 Oct 06  [8]
2600:0000::/12        ARIN           03 Oct 06  [9]
2610:0000::/23        ARIN           17 Nov 05
2620:0000::/23        ARIN           12 Sep 06
2800:0000::/12        LACNIC         03 Oct 06  [7]
2A00:0000::/12        RIPE NCC       03 Oct 06  [5]
2C00:0000::/12        AfriNIC        03 Oct 06
```

Networks 2008 - IPv6 tutorial

Production Addressing Scheme (5)

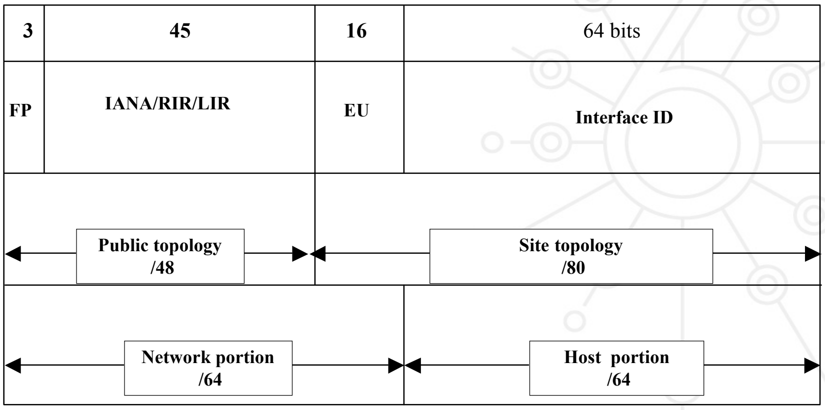

RIR Allocation Policies

AfriNIC:
http://www.afrinic.net/IPv6/index.htm
http://www.afrinic.net/docs/policies/afpol-v6200407-000.htm *
APNIC:
http://www.apnic.org/docs/index.html
http://www.apnic.org/policy/ipv6-address-policy.html *
ARIN:
http://www.arin.net/policy/index.html
http://www.arin.net/policy/nrpm.html#ipv6 *
LACNIC:
http://lacnic.net/sp/politicas/
http://lacnic.net/sp/politicas/ipv6.html *
RIPE-NCC:
http://www.ripe.net/ripe/docs/ipv6.html
http://www.ripe.net/ripe/docs/ipv6policy.html *

- *describes policies for the allocation and assignment of globally unique IPv6 address space

6deploy.org

RIR Allocation Statistics

AfriNIC:
- http://www.afrinic.net/statistics/index.htm

APNIC:
- http://www.apnic.org/info/reports/index.html

ARIN:
- http://www.arin.net/statistics/index.html

LACNIC:
- http://lacnic.org/sp/est.html

RIPE-NCC:
- http://www.ripe.net/info/stats/index.html

See http://www.ripe.net/rs/ipv6/stats/

Networks 2008 - IPv6 tutorial

IPv6 associated Protocols

Networks 2008 - IPv6 tutorial Budapest, 28 Sept 2008

6deploy.org

New Protocols

New features specified in IPv6 Protocol (RFC 2460 DS)

Neighbor Discovery (ND) (RFC 2461 DS)

Auto-configuration :

- Stateless Address Auto-configuration (RFC 2462 DS)
- DHCPv6: Dynamic Host Configuration Protocol for IPv6 (RFC 3315 PS)
- Path MTU discovery (pMTU) (RFC 1981 PS)

Networks 2008 - IPv6 tutorial

6deploy.org

New Protocols (2)

MLD (Multicast Listener Discovery) (RFC 2710 PS)

- Multicast group management over an IPv6 link
- Based on IGMPv2
- MLDv2 (equivalent to IGMPv3 in IPv4)

ICMPv6 (RFC 2463 DS) "Super" Protocol that :

- Covers ICMP (v4) features (Error control, Administration, …)
- Transports ND messages
- Transports MLD messages (Queries, Reports, …)

Networks 2008 - IPv6 tutorial

Neighbor Discovery

- **IPv6 nodes which share the same physical medium (link) use Neighbor Discovery (NDP) to:**

  discover their mutual presence
  determine link-layer addresses of their neighbors
  find routers
  maintain neighbors' reachability information (NUD)
  not directly applicable to NBMA (Non Broadcast Multi Access) networks ND uses multicast for certain services.

6deploy.org

Neighbor Discovery (2)

Protocol features:

- Router discovery
- Prefix(es) discovery
- Parameters discovery (link MTU, Max Hop Limit, ...)
- Address auto-configuration
- Address resolution
- Next Hop determination
- Neighbor Unreachability Detection
- Duplicate Address Detection
- Redirect

Networks 2008 - IPv6 tutorial

Neighbor Discovery (3): Comparison with IPv4

6deploy.org

It is the synthesis of:

- ARP
- R-Disc
- ICMP redirect
- ...

Networks 2008 - IPv6 tutorial

6deploy 6deploy.org

Neighbor Discovery (4)

ND specifies 5 types of ICMP packets :

- **Router Advertisement** (RA) :
  - periodic advertisement (of the availability of a router) which contains:
    - » list of prefixes used on the link (autoconf)
    - » a possible value for Max Hop Limit (TTL of IPv4)
    - » value of MTU
- **Router Solicitation** (RS) :
  - the host needs RA immediately (at boot time)

Networks 2008 - IPv6 tutorial

6deploy.org

Neighbor Discovery (5)

- **Neighbor Solicitation** (NS):
  - to determine the link-layer @ of a neighbor
  - or to check its reachability
  - also used to detect duplicate addresses (DAD)
- **Neighbor Advertisement** (NA):
  - answer to a NS packet
  - to advertise the change of physical address
- **Redirect** :
  - Used by a router to inform a host of a better route to a given destination

Networks 2008 - IPv6 tutorial

6deploy.org

Address Resolution

Find the mapping:

Destination IP @ ➔ Link-Layer (MAC) @

Recalling IPv4 & ARP

- ARP Request is broadcasted
  - e.g. ethernet @: FF-FF-FF-FF-FF-FF
  - Btw, it contains the Source's Link-local @
- ARP Reply is sent in unicast to the Source
  - It contains the Destination's Link-local @

Networks 2008 - IPv6 tutorial

Address Resolution (2) IPv6 with Neighbor Discovery

6deploy.org

Every IPv6 node MUST join 2 special multicast groups for each network interface:

- All-nodes multicast group: **ff02::1**
- Solicited-node multicast group

Concatenation **of the prefix FF02::1:FF00:0/104 with the last 24 bits of the IPv6 address**

Dst IPv6 **@: 2001:0660:010a:4002:4421:21FF:FE**24:87c1

↓

Sol. Mcast **@:** FF02:0000:0000:0000:0000:0001:FF24:87c1

↓

ethernet: 33-33-FF-24-87-c1

Networks 2008 - IPv6 tutorial

Address Resolution (3) IPv6 with Neighbor Discovery

6deploy.org

- H1 knows IP address of H2 (IP2) and wants to know its MAC address (MAC2)
- H1 builds the sollicited multicast address of IP2 : Multi (IP2)
- H1 sends « Neighbor sollicitation » message to this sollicited multicast IPv6 address
- At **link level**, NS packet is sent to a **multicast address** instead of broacast

Networks 2008 - IPv6 tutorial

Address Resolution (4) IPv6 with Neighbor Discovery

- Ethernet manages multicast – Not always implemented
- Ethernet frame is often broadcasted on the link
- Only H2 is destination of the ethernet frame and sends the « Neighbor Sollicitation » packet to its IPv6 stack
- H2 replies sending a unicast « Neighbor Advertisement » message to H1. This message contains the link layer address of H2.

Networks 2008 - IPv6 tutorial

Path MTU discovery (RFC 1981)

Derived from RFC 1191, (IPv4 version of the protocol)

Path **: set of links followed by an IPv6 packet between source and destination**

link MTU **: maximum packet length (bytes) that can be transmitted on a given link without fragmentation**

Path MTU **(or pMTU) = min { link MTUs } for a given path**

Path MTU Discovery = automatic pMTU discovery for a given path

Networks 2008 - IPv6 tutorial

Path MTU discovery (2)

Protocol operation

- makes assumption that pMTU = link MTU to reach a neighbor (first hop)
- if there is an intermediate router such that link MTU < pMTU ➔ it sends an ICMPv6 message: "Packet size Too Large"
- source reduces pMTU by using information found in the ICMPv6 message

=> Intermediate equipments aren't allowed to perform packet fragmentation

Networks 2008 - IPv6 tutorial

Auto-configuration

Hosts should be plug & play

Use ICMPv6 messages (Neighbor Discovery)

When booting, the host asks for network parameters:

- IPv6 prefix(es)
- default router address(es)
- hop limit
- (link local) MTU
- ...

Networks 2008 - IPv6 tutorial

Auto-configuration (continued)

Only routers have to be manually configured

- but work on **prefix delegation** is in progress
  (draft-ietf-ipv6-prefix-delegation-requirement-01.txt)

Hosts can get automatically an IPv6 address

- BUT it is not automatically registered in the DNS
- If the address is always the same: may be manually registered

NEED for DNS Dynamic Update

(RFC 2136 PS and RFC 3007 PS) for IPv6

- Security issues …

Networks 2008 - IPv6 tutorial

Stateless auto-configuration

IPv6 Stateless Address Auto-configuration

- RFC 2462 DS
- Does not apply to routers

Allows a host to create a global IPv6 @ from:

- Its interface identifier = EUI64(MAC @)
- router advertisements coming from router(s) on the link

=> GA = concat (RA, EUI64)

Networks 2008 - IPv6 tutorial

6deploy.org

Auto-configuration example

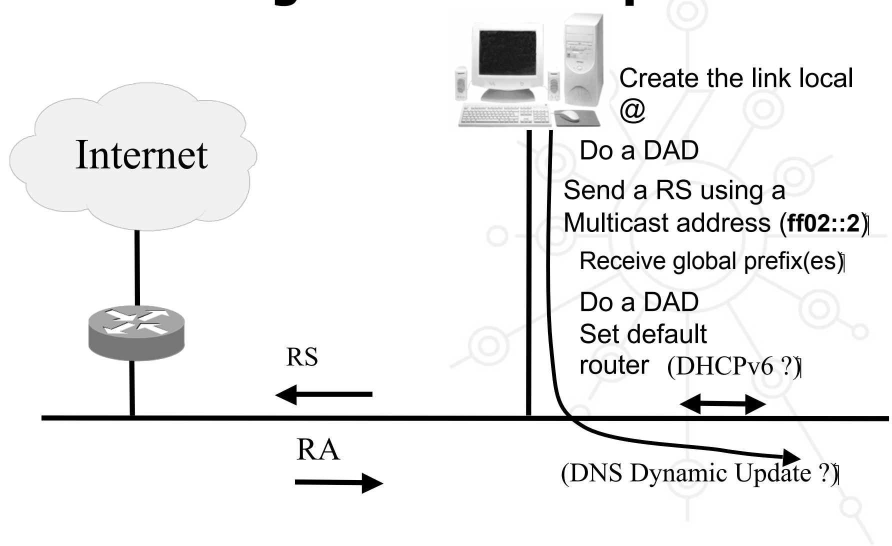

Networks 2008 - IPv6 tutorial

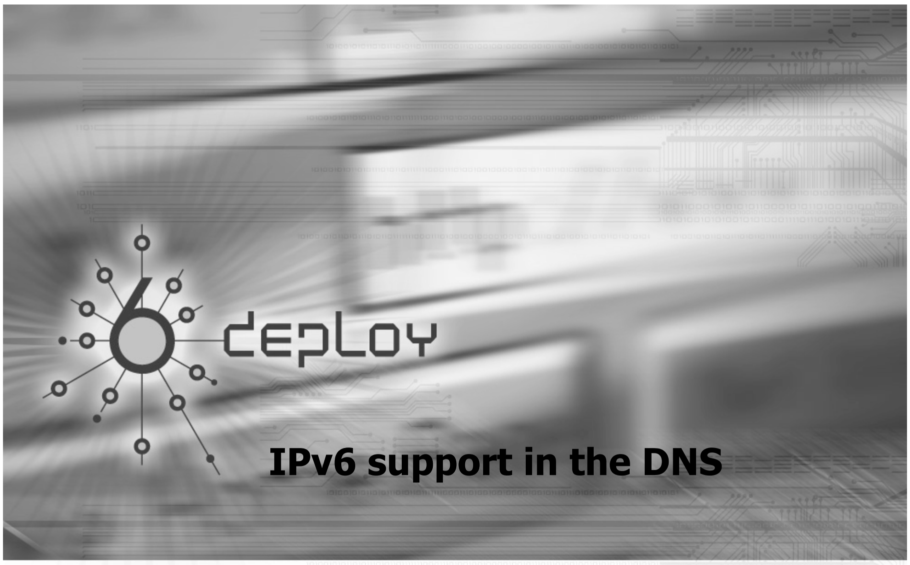

Networks 2008 - IPv6 tutorial   Budapest, 28 Sept 2008

DNS Extensions for IPv6

6deploy.org

- RFC 1886 (PS) → RFC 3596 (DS) (upon successful interoperability tests)
- **AAAA** (RFC 3596): forward lookup ('Name → IPv6 Address'):
  - Equivalent to '**A**' record
  - Example:

```
ns3.nic.fr.   IN   A      192.134.0.49
              IN   AAAA   2001:660:3006:1::1:1
```

- **PTR** : reverse lookup ('IPv6 Address → Name'):
  - Reverse tree equivalent to in-addr.arpa
    - Nibble (4 bits) boundary
    - New tree: ip6.arpa (RFC 3596), used
    - Former tree: ip6.int (RFC 1886), obsolote
  - Example:

```
$ORIGIN 1.0.0.0.6.0.0.3.0.6.6.0.1.0.0.2.ip6.{int,arpa}.
1.0.0.0.1.0.0.0.0.0.0.0.0.0.0.0   PTR   ns3.nic.fr.
```

Networks 2008 - IPv6 tutorial

6deploy.org

IPv6 DNS and root servers

DNS root servers are critical resources!

13 roots « around » the world (#10 in the US)

As of 04/02/2008, 6 of 13 root servers are IPv6 enabled and reachable via IPv6 networks (A, F, H, J, K, M).

IPv6 DNS operation

Recent BIND, NSD, PowerDNS - OK

Microsoft Windows XP default resolver only queries over IPv4 transport:

- Install BIND 9 for Windows XP and uses BINDs resolver

The target today IS NOT the transition from an IPv4-only to an IPv6-only environment

- Start by testing DNSv6 on a small network and get your own conclusion that DNSv6 is harmless
- Deploy DNSv6 in an incremental fashion on existing networks
- DO NOT BREAK something that works fine (production IPv4 DNS)!

Networks 2008 - IPv6 tutorial

Routing Protocols

Internal and External Routing

Networks 2008 - IPv6 tutorial   Budapest, 28 Sept 2008

6deploy.org

Agenda

Internal Routing

- RIPng
- IS-IS
- OSPFv3

External Routing

- Multiprotocol BGP

Networks 2008 - IPv6 tutorial

6deploy.org

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

Networks 2008 - IPv6 tutorial

6deploy.org

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

Networks 2008 - IPv6 tutorial

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

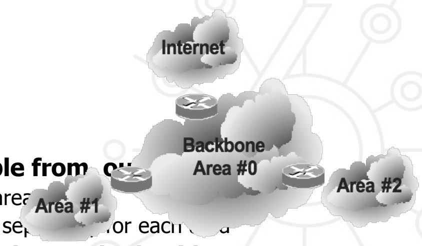

Topology of an area is invisible from ou
- LSA flooding is bounded by are
- SPF calculation is performed se or each
All areas must have a connection to the backbone

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
- Distributes IPv6 prefixes
- New LSA types
- Uses the Multicast address
  - ALLSPFRouters (FF02::5)
  - ALLDRouters (FF02::6)

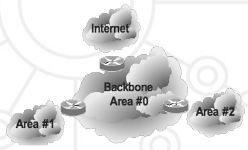

Multiprocol BGP

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

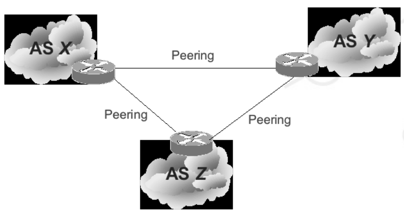

6deploy.org

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

Networks 2008 - IPv6 tutorial

6deploy.org

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)

Networks 2008 - IPv6 tutorial

6deploy.org

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

Networks 2008 - IPv6 tutorial

Campus IPv6 deployment - Campus Address allocation, Topology Issues
János Mohácsi
NIIF/HUNGARNET

Networks 2008 - IPv6 tutorial Budapest, 28 Sept 2008

Outline

Campus deployment strategy
Campus IPv6 address allocation
Campus deployment topology - options
Campus services

Service provider deployment considerations

Networks 2008 - IPv6 tutorial

Outline

Campus deployment strategy
Campus IPv6 address allocation
Campus deployment topology - options
Campus services

Service provider deployment considerations

Networks 2008 - IPv6 tutorial

Various Campus transition approaches

IPv4 will be used for years after IPv6 deployment

Dual Stack

- Servers/clients speaking both protocols
- Application/service can select either protocol to use

Tunneling ("connecting IPv6 clouds")

- IPv6 packet is data payload of IPv4 packet/or MPLS frames

Translation methods ("IPv4<->IPv6 services")

- Layer 3: Rewriting IP header information (NAT-PT)
- Layer 4: Rewriting TCP headers
- Layer 7: Application layer gateways (ALGs)

Benefits of dual-stack deployment

By deploying dual-stack, you can test IPv6-only devices/services without disrupting IPv4 connectivity

Dual stack IPv6 + IPv4 NAT: legacy IPv4 applications (email, www) can be used next to new IPv6 applications (p2p, home networking, …)

- IPv6 offers the next generation of applications

Campus deployment plan /1

1. **Obtain global IPv6 address space from your ISP**
   - NRENs usually has a /32 prefix from RIPE NCC/RIRs
   - A universities/customers will get a /48 prefix from NRENs/LIRs
2. **Obtain external connectivity**
   - You can do dual-stack connectivity
   - Many universities will use tunnel to to get IPv6 service
     - in this case be sure that nobody can abuse your tunnel – use filtering

Campus deployment plan /2

3. **Internal deployment**
   - Determine an IPv6 firewall/security policy
   - Develop an IPv6 address plan for your site
   - Determine address management policy (RA/DHCPv6?)
   - Migrate to dual-stack infrastructure on the wire
     - Network links become IPv6 enabled
   - Enable IPv6 services and applications
     - Starting with DNS
   - Enable IPv6 on host systems (Linux, WinXP, …)
   - Enable management and monitoring tools

Outline

- **Campus deployment strategy**
- **Campus IPv6 address allocation**
- **Campus deployment topology - options**
- **Campus services**

- **Service provider deployment considerations**

Campus Addressing

- **Most sites will receive /48 assignments:**

Network address (48 bits)	16bits	EUI host address (64 bits)

- **16 bits left for subnetting - what to do with them?**

6deploy.org

Campus Address - site level subnetting

1. Sequentially, e.g.

- 0000
- 0001
- …
- FFFF

- 16 bits = 65535 subnets

Networks 2008 - IPv6 tutorial

6deploy.org

Campus Addressing - site level subnetting

2. Following existing IPv4:

- Subnets or combinations of nets & subnets, or VLANs, etc., e.g.
- IPv4 subnets:
  - `152.66.`**`60`**`.0/24        003c`
  - `152.66.`**`91`**`.0/24        005b`
  - `152.66.`**`156`**`.0/24       009c`
- VLANs:
  - `VLAN id 100          0100 (w/o conversion) or 0064 (w conversion)`

Networks 2008 - IPv6 tutorial

Campus Addressing - site level subnetting

3. Topological/aggregating

reflecting wiring plants, supernets, large broadcast domains, etc.

- Main library = 0010/60
  - Floor in library = 001a/64
- Computing center = 0200/56
  - Student servers = 02c0/64
- Medical school = c000/52
- and so on. . .

New Things to Think About

You can use "all 0s" and "all 1s"! (0000, ffff)

You're not limited to 254 hosts per subnet!

- Switch-rich LANs allow for larger broadcast domains (with tiny collision domains), perhaps thousands of hosts/LAN…

No "secondary address" (though >1 address/interface)

No tiny subnets either (no /30, /31, /32)—plan for what you need for backbone blocks, loopbacks, etc.

You should use /64 per links - Especially if you plan to use autoconfiguration!

New Things to Think About /2

Every /64 subnet has far more than enough addresses to contain all of the computers on the planet, and with a /48 you have 65536 of those subnets - use this power wisely!

With so many subnets your IGP may end up carrying thousands of routes - consider internal topology and aggregation to avoid future problems.

Networks 2008 - IPv6 tutorial

New Things to Think About /3

Renumbering will likely be a fact of life. Although v6 does make it easier, it still isn't pretty. . .

- Avoid using numeric addresses at all costs
- Avoid hard-configured addresses on hosts except for servers (this is very important for DNS servers) – use the feature that you can assign more than one IPv6 address to an interface (IPv6 alias address for servers)
- Anticipate that changing ISPs will mean renumbering

Networks 2008 - IPv6 tutorial

New Things to Think About /4

Recap from EUI-64:

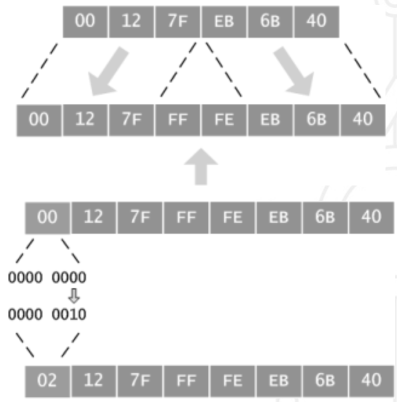

- The motivation for inverting the 'u' bit when forming the interface identifier is to make it easy for system administrators to hand configure local scope identifiers. This is expected to be case for serial links, tunnel end-points and servers, etc. simply ::1, ::2, etc

Outline

Campus deployment strategy
Campus IPv6 address allocation
Campus deployment topology - options
Campus services

Service provider deployment considerations

6deploy.org

IPv6 deployment options

- The simplest - deploy dual stack network environment
- If the hosts/services are not dual stack enabled - does not break anything - this tend to be a false assumption (Windows Vista, Mac OS X shipped with IPv6 enabled)
- If the L3 devices cannot cope with IPv6 or administrators are not favor of upgrading the router
  - Add additional IPv6 capable L3 device

Networks 2008 - IPv6 tutorial

6deploy.org

Layer-2 Campus - 1 Switch

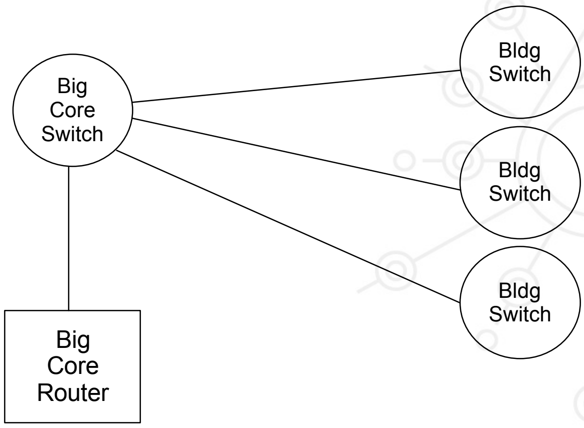

Networks 2008 - IPv6 tutorial

Layer-2 Campus - 1 Switch

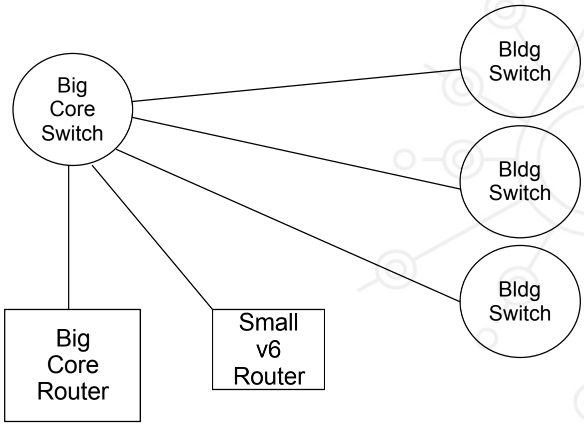

Networks 2008 - IPv6 tutorial

Layer-2 Campus - Redundant Switches

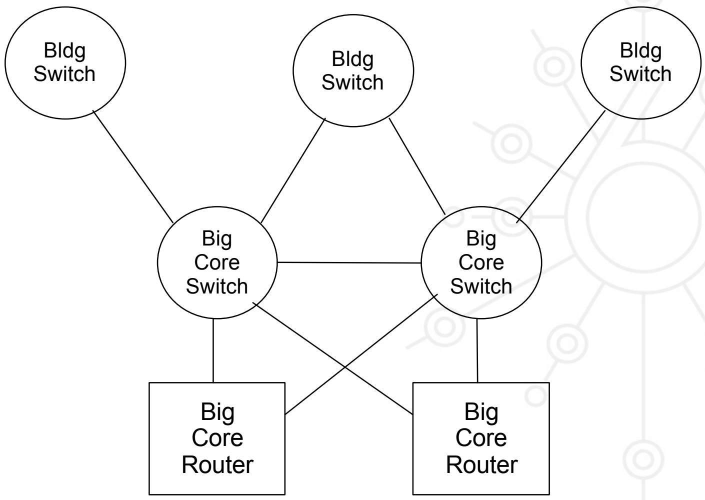

Networks 2008 - IPv6 tutorial

Layer-2 Campus Redundant Switches

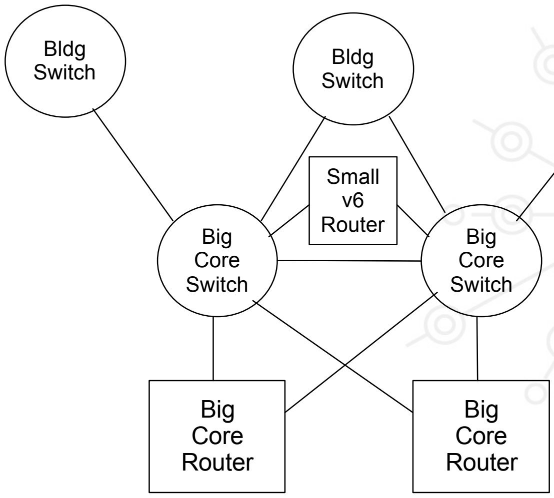

Layer-3 Campus

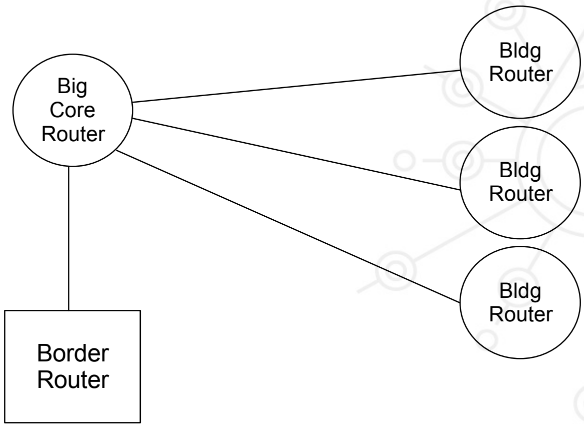

Layer-3 Campus

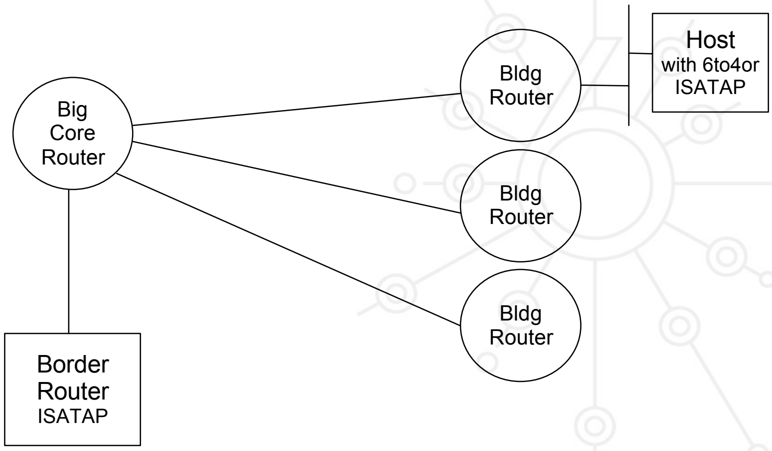

Layer-3 Campus

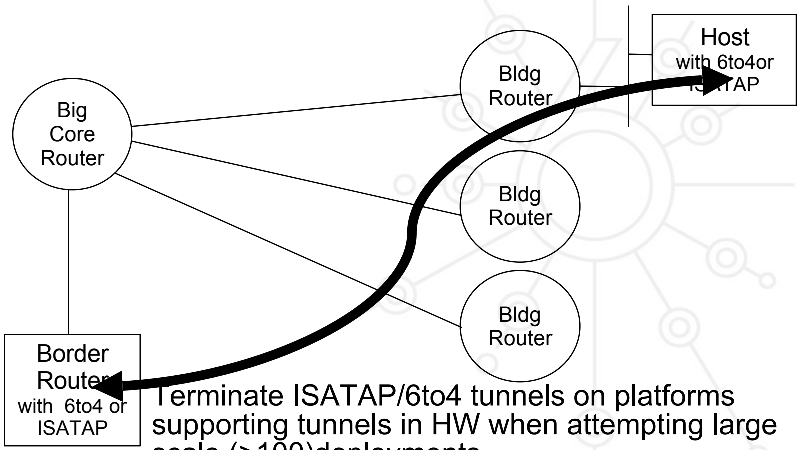

Terminate ISATAP/6to4 tunnels on platforms supporting tunnels in HW when attempting large scale (>100)deployments

Edge Router Options

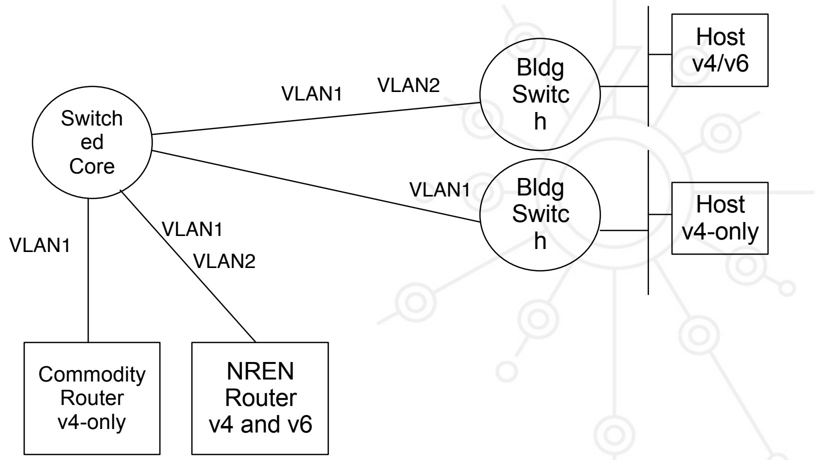

Routing Protocols

iBGP and IGP (IS-IS/OSPFv3)

- IPv6 iBGP sessions in parallel with IPv4
- You need IPv4 router-id for IPv6 BGP peering

Static Routing

- all the obvious scaling problems, but works OK to get started, especially using a trunked v6 VLAN.

OSPFv3 is might be good

- It will run in a ships-in-the-night mode relative to OSPFv2 for IPV4 - neither will know about the other.

Outline

Campus deployment strategy
Campus IPv6 address allocation
Campus deployment topology - options
Campus services

Service provider deployment considerations

IPv6 load balancing

- Server clusters
  - Opensource solution: *BSD pf
  - Commercial platforms: Veritas Cluster Server, BigIron F5, Windows Server 2008 - Network Load Balancer
- First-Hop Redundancy:
  - HSRPv6 (Cisco only)
  - VRRPv6 - standardisation at IETF
  - NUD - see next slide
- Traffic loadbalancing
  - Multilink PPP - supported if multilink PPP supported
  - Equal-Cost Multi-Path routing - if IPv6 routing supported…
  - Ethernet Link Aggregations - L2 solution

Implementing default gateway redundancy

If HSRP,GLBP or VRRP for IPv6 are not available

NUD can be used for a good HA at the first-hop (today this only applies to the Campus/DC…HSRP is available on routers)

- (config-if)#ipv6 nd reachable-time 5000

Hosts use NUD "reachable time" to cycle to next known default gateway (30 seconds by default)

Default Gateway . . . . : 10.121.10.1

fe80::211:bcff:fec0:d000%4

fe80::211:bcff:fec0:c800%4

Reachable Time : 6s

Base Reachable Time : 5s

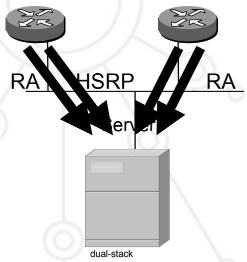

Management and monitoring

- Device configuration and monitoring -SNMP
- Statistical monitoring e.g. Cricket/MRTG
- Service monitoring - Nagios
- Intrusion detection (IDS)
- See more later

How to enable IPv6 services?

Add v6 testing service for different name first:

- service.v6.fqdn or service6.fqdn with AAAA + reverse PTR entry.
- Test it

Add v6 service under the same name:

- service.fqdn with A +AAAA and two PTR.

How to enable IPv6 services if you don’t have IPv6 capable server?

Use proxy (more exactly reverse-proxy) server

- Apache2 proxy is a very good one

Use netcat

- Kind of hack ☺

Other proxies

Proxy solutions

- **Proxy**
  - Squid (http://devel.squid-cache.org/projects.html)
- **Web Cache**
  - NetCache C1300, C2300, C3300. BlueCoat SG
  - WCCP does not have IPv6 support in CISCO yet

Networks 2008 - IPv6 tutorial

Apache2 reverse proxy

Configuration is very easy:

```
ProxyRequests Off
ProxyPass / http://ipv4address
ProxyPassReverse / http://ipv4address
ProxyPreserveHost On
```

Networks 2008 - IPv6 tutorial

6deploy.org

Reverse proxy advantages & disadvantages

Advantage:

- Fast implementation, instantly provide web service over IPv6
- No modifications required in a production web server environment
- Allow for timely upgrading of systems
- Scalable mechanism: a central proxy can support many web sites

Disadvantage:

- Significant administrative overhead for large scale deployment
- May break advanced authentication and access control schemes
- Breaks statistics: all IPv6 requests seem to be coming from the same address (may be fixed with filtering and concatenation of logs)
- Not a long term solution overall, native IPv6 support is readily available in related applications and should be preferred whenever possible

Networks 2008 - IPv6 tutorial

6deploy.org

DHCP (1)

IPv6 has stateless address autoconfiguration but DHCPv6 (RFC 3315) is available too

DHCPv6 can be used both for assigning addresses and providing other information like nameserver, ntpserver etc

If not using DHCPv6 for addresses, no state is required on server side and only part of the protocol is needed. This is called Stateless DHCPv6 (RFC 3736)

Some server and client implementations only do Stateless DHCPv6 while others do the full DHCP protocol

The two main approaches are

- Stateless address autoconfiguration with stateless DHCPv6 for other information
- Using DHCPv6 for both addresses and other information to obtain better control of address assignment

Networks 2008 - IPv6 tutorial

DHCP (2)

One possible problem for DHCP is that DHCPv4 only provides IPv4 information (addresses for servers etc) while DHCPv6 only provides IPv6 information. Should a dual-stack host run both or only one (which one)?

Several vendors working on DHCP but only a few implementations available at the moment

- DHCPv6 http://dhcpv6.sourceforge.net/
- dibbler http://klub.com.pl/dhcpv6/
- NEC, Lucent etc. are working on their own implementations
- KAME – only stateless

Cisco routers have a built-in stateless server that provides basic things like nameserver and domain name (also SIP server options in image I checked).

DHCP can also be used between routers for prefix delegation (RFC 3633). There are several implementations. E.g. Cisco routers can act as both client and server

Networks 2008 - IPv6 tutorial

Remote access via IPv6

Use native connectivity –

- Rather easy if you are operating dial-in pool or you are an ADSL service provider

Use 6to4 if you have global IPv4 address

- Good 6to4 relay connectivity is a must

Use tunnelbroker service – rather suboptimal

Use Teredo/softwire if you have NAT our multiple level of NATs

Use OpenVPN

Networks 2008 - IPv6 tutorial

Remote access via IPv6 - PPP

- **The dial-up connection uses a modem and the PSTN service in order to get connection to remote devices.**
  - Most cases is used the PPP (Point-to-Point Protocol), which gives a standard method to transport the datagrams of several protocols over point-to-point links (RFC1661) - PPP has been updated to support the transport of IPv6 datagrams (RFC5072)

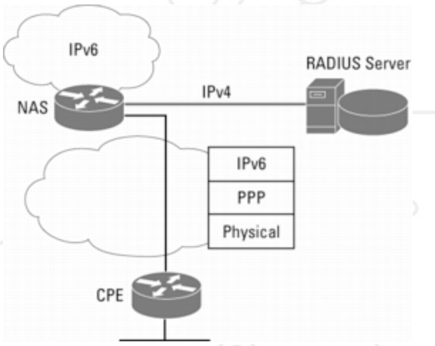

Networks 2008 - IPv6 tutorial

PPP and IPv6

- PPP protocol has three main parts
  - Definition of the encapsulation method of the IPv6 datagrams over the over the point- to-point link (IP6CP)
  - LCP (Link Control Protocol) use to establish, configure and test the connection at link layer
  - NCP (Network Control Protocol) use to establish and configure the connection at network layer
- IPv6 operation:
  - negotiates one link local address (fe80::/64) between the end points or peers
  - Could negotiate datagram compression via IPV6CP (IPv6 Control Protocol)
  - PPP does not give global IPv6 addresses but link local - The global IPv6 addresses must be configured by other means
    - Manual configuration
    - Autoconfiguration (RA)
    - DHCPv6

Networks 2008 - IPv6 tutorial

PPP and IPv6 - implementations

- Routers:
  - Cisco
  - Juniper
- Hosts:
  - Windows Vista and Microsoft Windows Server 2008
  - Linux, *BSD, Solaris
- Opensource:

  http://sourceforge.net/projects/pppcbcp
  http://freshmeat.net/projects/pppd

Remote Access with IPSEC – or other VPNs

IPv6-in-IPv4 Tunnel Example

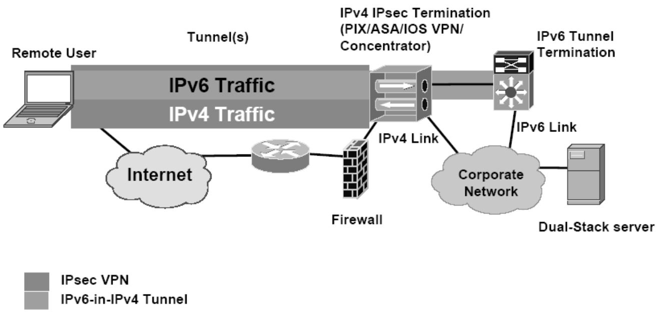

Outline

Campus deployment strategy
Campus IPv6 address allocation
Campus deployment topology - options
Campus services

Service provider deployment considerations

Outline of NRENs/ISP IPv6 deployment

1. Obtain IPv6 address space
2. Plan the addressing
3. Plan the routing
4. Test in a small case
5. Deploy IPv6 (incrementally – dual-stack/6PE)
6. Enable IPv6 services

Getting IPv6 prefix for LIRs/ISPs

Global IPv6 RIR rules

- http://www.ripe.net/ripe/docs/ipv6.html
- simple rules for LIRs
- IPv6 service should be provided
- detailed plan
- Usually /32 allocation

Establishing global rules was not easy.

- Different structure in different RIR regions: ISP, NIRs/LIRs, LIRs

What about IX? – slightly different rules

- Infrasstructure addresses
- Routable /48 address

IPv6 RIPE entries/1

```
whois -h whois.ripe.net 2001:0738::

inet6num:      2001:0738::/32
netname:       HU-HUNGARNET-20010717
descr:         Hungarnet IPv6 address block
               Hungarian Research & Educational Network
               Budapest, Hungary
country:       HU
mnt-by:        RIPE-NCC-HM-MNT        ←New  mandatory
mnt-lower:     NIIF6-MNT              ←New mandatory
status:        ALLOCATED-BY-RIR       ←New
```

IPv6 RIPE entries/2

possible values of STATUS field

- ALLOCATED-BY-RIR – Allocated address space by RIR to LIR.
- ALLOCATED-BY-LIR – Allocated address space by LIR to smaller registries/institutions
- ASSIGNED – Assigned to end-users

RPSLng ready for testing

Reverse delegation is strongly recommended

Questions?

6DEPLOY Project Web Site:

http://www.6deploy.eu

mohacsi@niif.hu