



deploy

Campus IPv6 deployment



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

- János Mohácsi, NIIF/HUNGARNET - Hungary

Contributors

- Jérôme Durand, Renater, France
- Tim Chown, University of Southampton, Great-Britain
- B. Tuy, Renater, France

Warning ...

This module is under work (it's evolving still rapidly ...)

- *here are ideas drawn from experienced people*
- *it's out of scope to recommend every one to do the same*
- *Every campus is specific and thinking what to do and how to do it beforehand is a must*

Good luck !

Outline

Campus deployment strategy

Campus IPv6 address allocation

Campus deployment topology - options

Campus services

Service provider deployment considerations

Outline

Campus deployment strategy

Campus IPv6 address allocation

Campus deployment topology - options

Campus services

Service provider deployment considerations



Various Campus transition approaches

*IPv4 will be used for years after IPv6 has been deployed
Then both versions of the IP protocol will have to coexist*

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 have a /32 prefix from RIPE NCC/RIRs
- Universities/customers will get a /48 prefix from NRENs/LIRs

2. Obtain external connectivity

- You can do dual-stack connectivity
- Many universities will use a tunnel 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
 - The IPv4 firewall/security policy is a good start
- Develop an IPv6 address plan for your site
- Determine an 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, Vista, Mac OS X...)
- 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:



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

Two main questions to answer:

⇒ **How many topologically different “zones” can be identified ?**

- Existing ones or new ones to be created for whatever (good) reason

⇒ **How many networks (subnets) are needed within these zones ?**

Example network « zones »

Zone description	Nb of subnets
Upstream interco and infrast	16
Administration services	4
Medical Sciences dept	32
Dept A	16
Dept B	16
...	

Campus Addressing - site level subnetting - methods -1

1. Sequentially, e.g.

- 0000
- 0001
- ...
- FFFF

- 16 bits = 65535 subnets

⇒ Reserve prefixes for further allocation

0020/60

0030/60

Subnet ID	Zone description
0000 / 60	BB Infrastructure
0010 / 60	Administration
0020 / 59	Medical Sciences dept
0040 / 60	Dept A
0050 / 60	Dept B
...	...

Campus Addressing - site level subnetting - methods 2

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 decimal/hex conversion)
 or 0064 (w dec/hex conversion)

⇒ *Best to start thinking about it*

Campus Addressing - site level subnetting - methods 3

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. . .



Example network - topological aggregation + sequential allocation

Zone description	Nb of subnets
Upstream interco and infrast	16
Administration services	4
Medical Sciences dept	32
Dept A	16
Dept B	16
...	

IPv6 subnet prefix allocations (ex.)

Subnet ID	Subnet prefix allocation	Description
0000 / 60		BB Infrastructure
	0000/64	Upstream interconnection
	0001/64	Campus architecture (DMZ)
	...	
	000B/64	Campus architecture
	...	
	000F	...
0010 / 60		Administration
	0010/64	Campus interco
	0011/64	Registration
	0012/64	Finance dept

IPv6 subnet prefix allocations ex. /2

Subnet ID	Subnet prefix allocation	Description
0020 / 60		Medical Sciences dept
	0020/64	Upstream interconnection
	0021/64	Nobel group
	...	
0030 / 60	Reserved	Medical Sciences dept
0040 / 60		Dept A
...		...

New Things to Think About

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

You’re not limited to the usual 254 hosts per subnet!

- LANs with lots of L2 switch 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!
- If you allocate global addressess interconnection links - not necessary in every case

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.

New Things to Think About /3

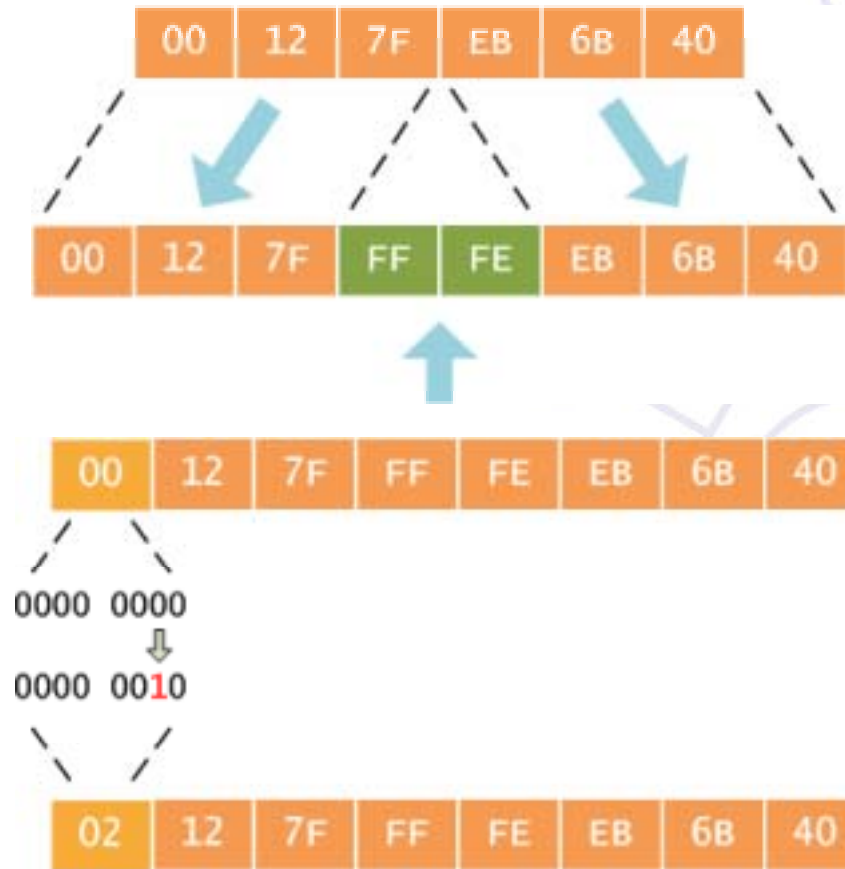
Renumbering will likely be a fact of life. Although v6 does make it easier, it still isn't pretty. . .

- 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
- An ISP change will impact the first 48 bits, you can keep the last 80 unchanged in every host/server's address.

**Address conservation usually not an issue
DHCPv6 might help**

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

Campus Addressing - address assignment

- Which address assignment to use?
 - Autoconfiguration - IEEE provides uniqueness
 - DHCPv6 - central management provides uniqueness
 - Manual - 7th bit of IID should be 0

Methods to manually assign addresses:

IID part	Description
0000:: <small>number</small>	Easy to remember allocations
0080:vvww:yyzz:XXXX/112	Automatically assigned to vv.ww.yy.zz IPv4 address: /112 belongs to a IPv4 host - good for service virtualisation

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 DHCPv6 is not used for address allocation, 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

- Some vendors don't implement yet a DHCPv6 client (MacOS X, ...)

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

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 integrations - several 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-WIDE DHCPv6 <http://sourceforge.net/projects/wide-dhcpv6/>
- ISC DHCPv6 <https://www.isc.org/software/dhcp>

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

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

Outline

Campus deployment strategy

Campus IPv6 address allocation

Campus deployment topology - options

Campus services

Service provider deployment considerations



IPv6 deployment options

The simplest

- deploy dual stack network environment

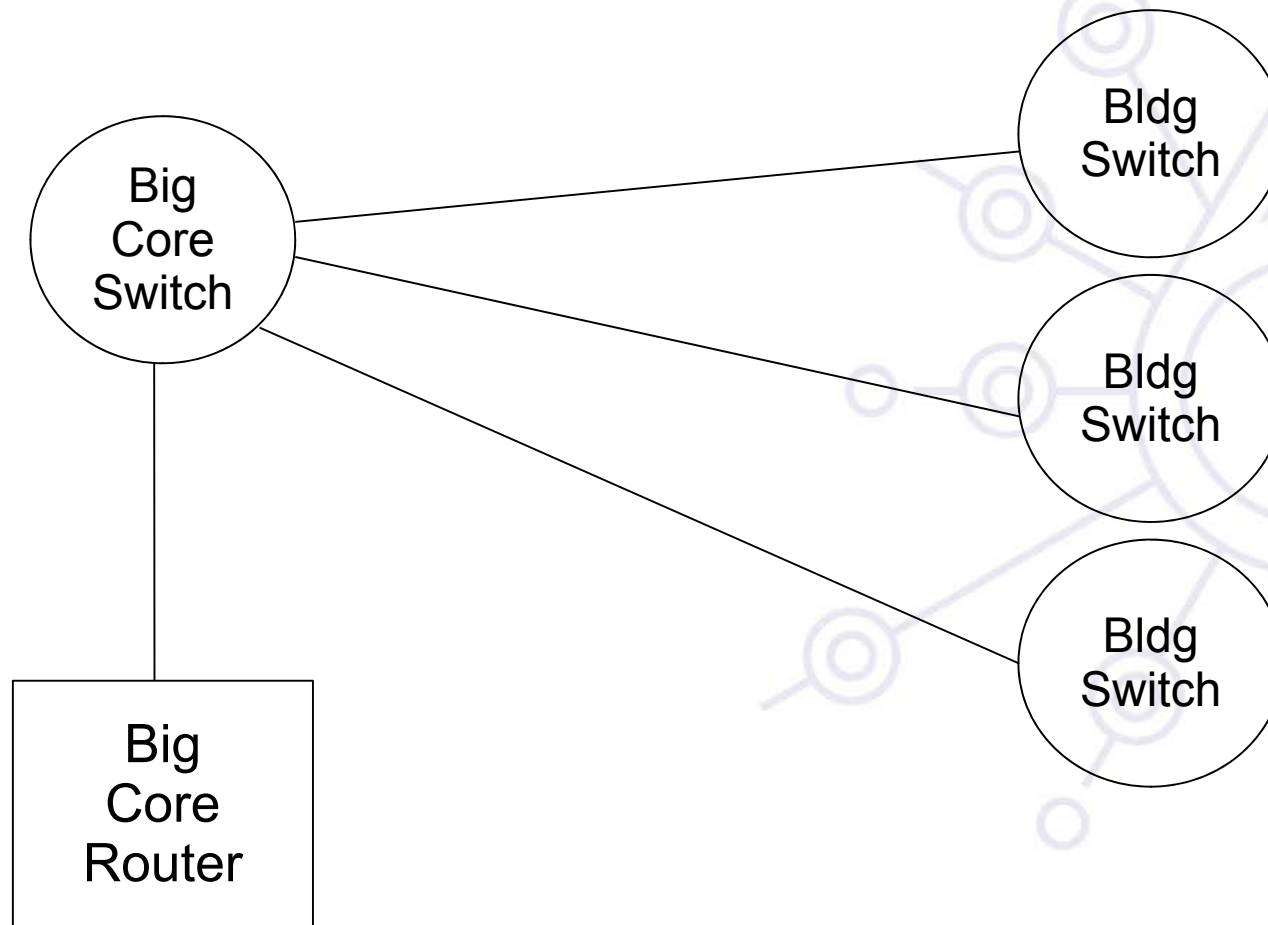
If the hosts/services are not dual stack enabled

- It does not break anything
- this tends to be a false assumption (Windows Vista, Mac OS X shipped with IPv6 enabled)

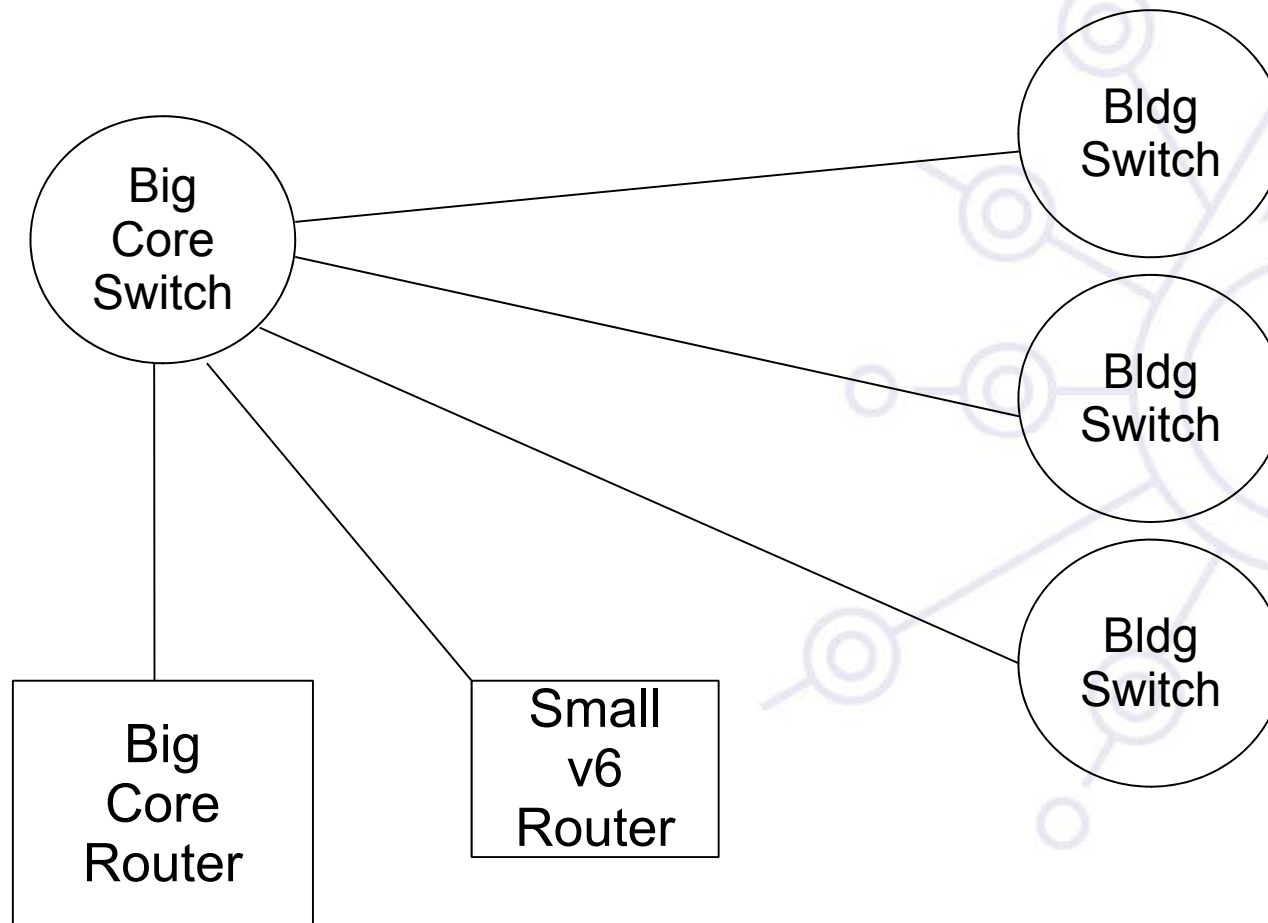
If the L3 devices cannot cope with IPv6 or administrators are not in favor of upgrading the router

- Add additional IPv6 capable L3 device(s)
- Investment money is usually a problem, but you can do some engineering with simple (low cost) PCs.

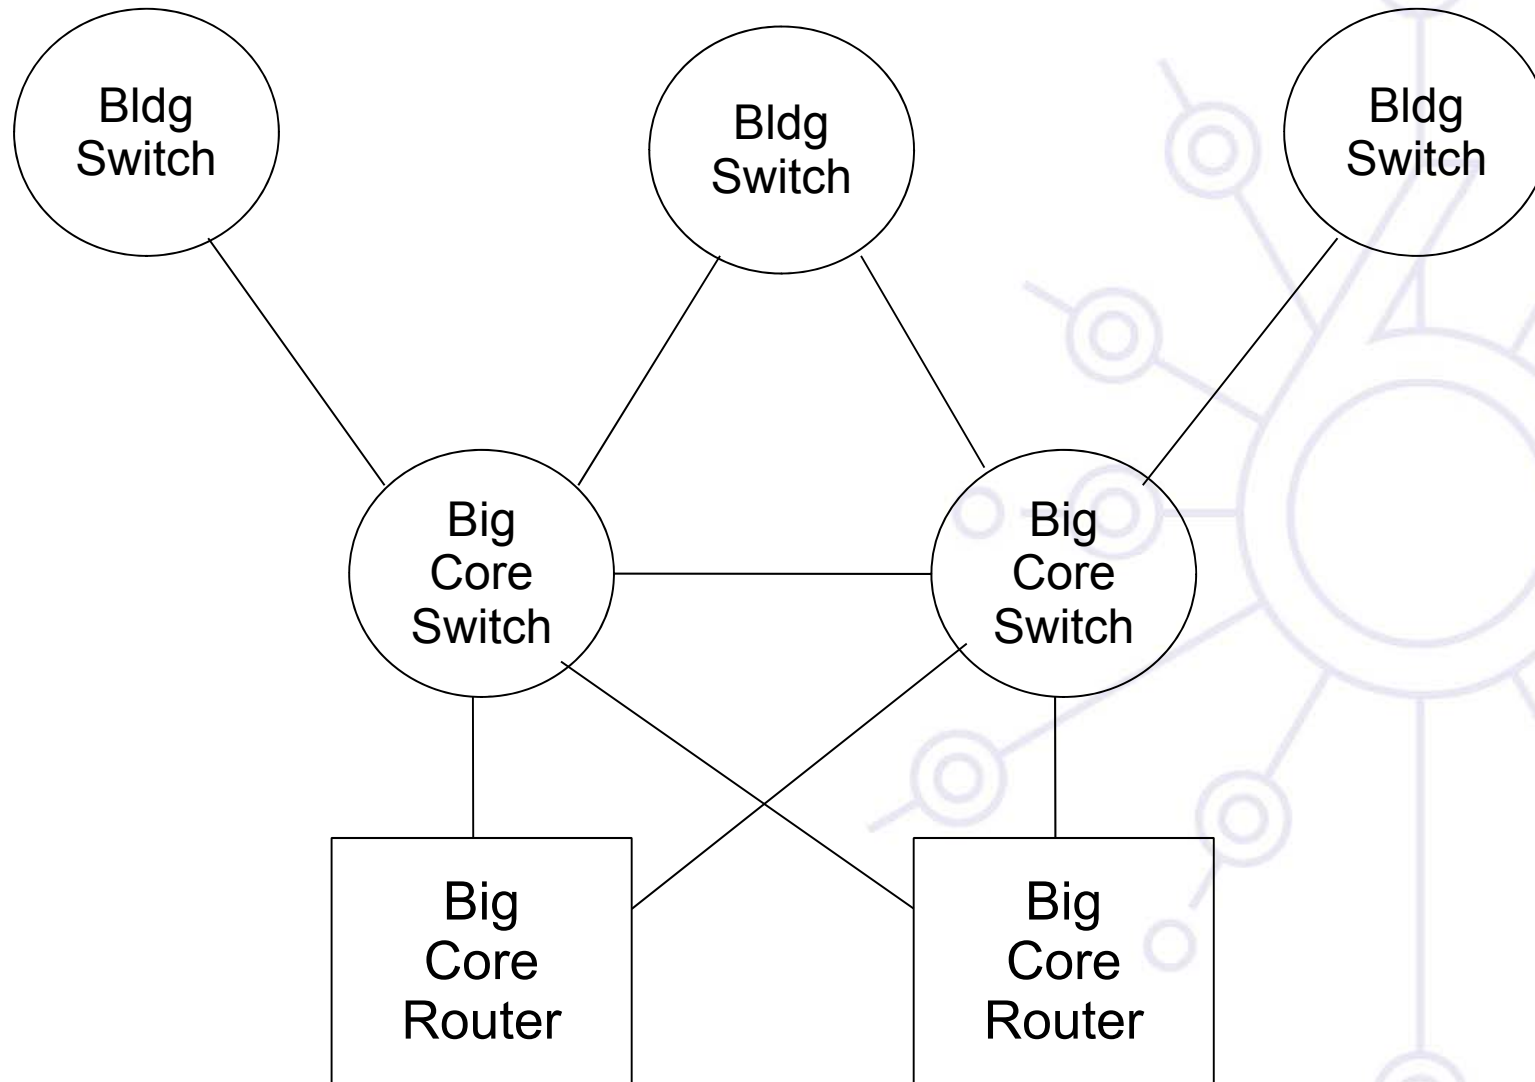
Layer-2 Campus - 1 Switch




Layer-2 Campus - 1 Switch



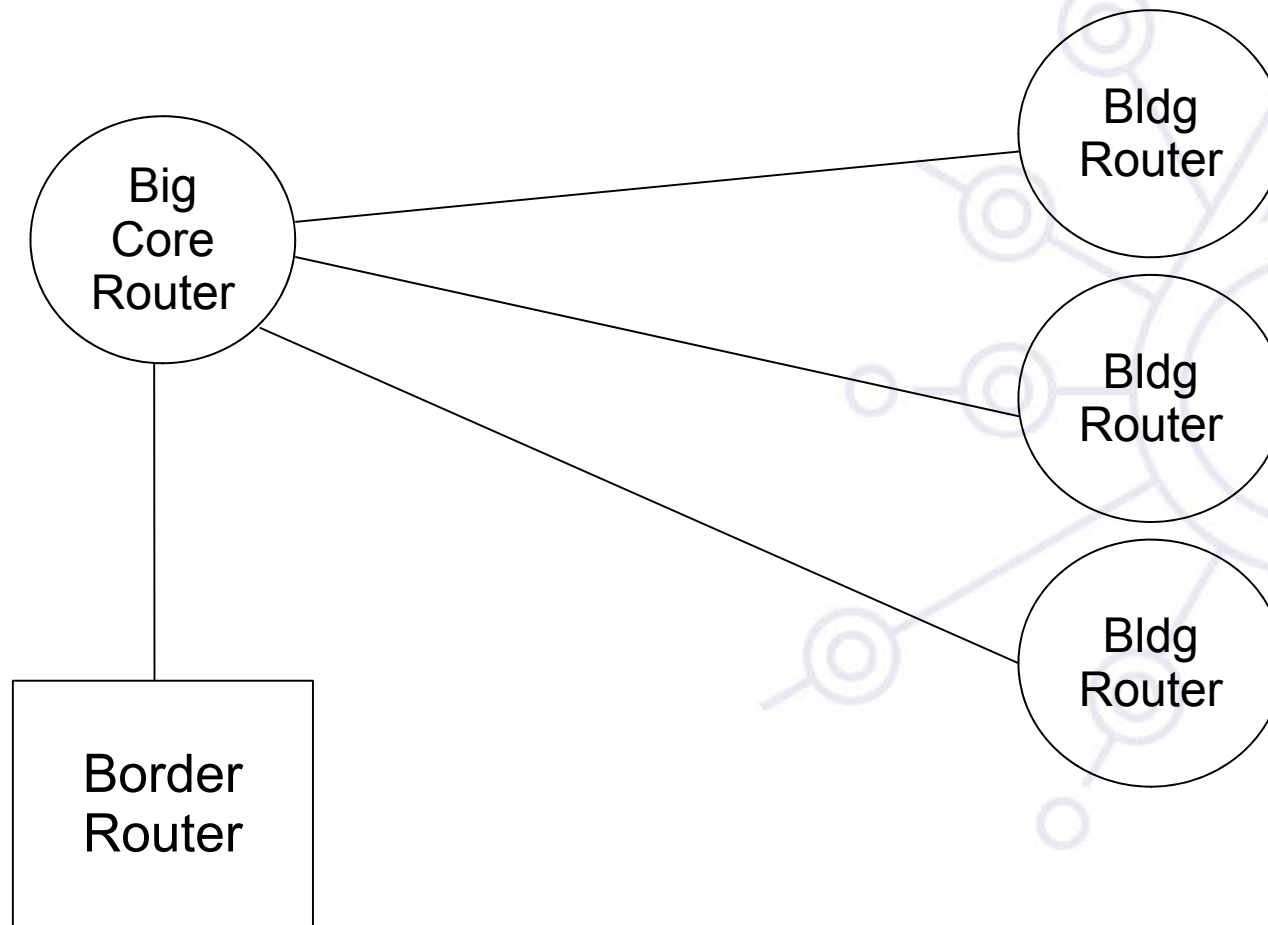
Layer-2 Campus - Redundant Switches



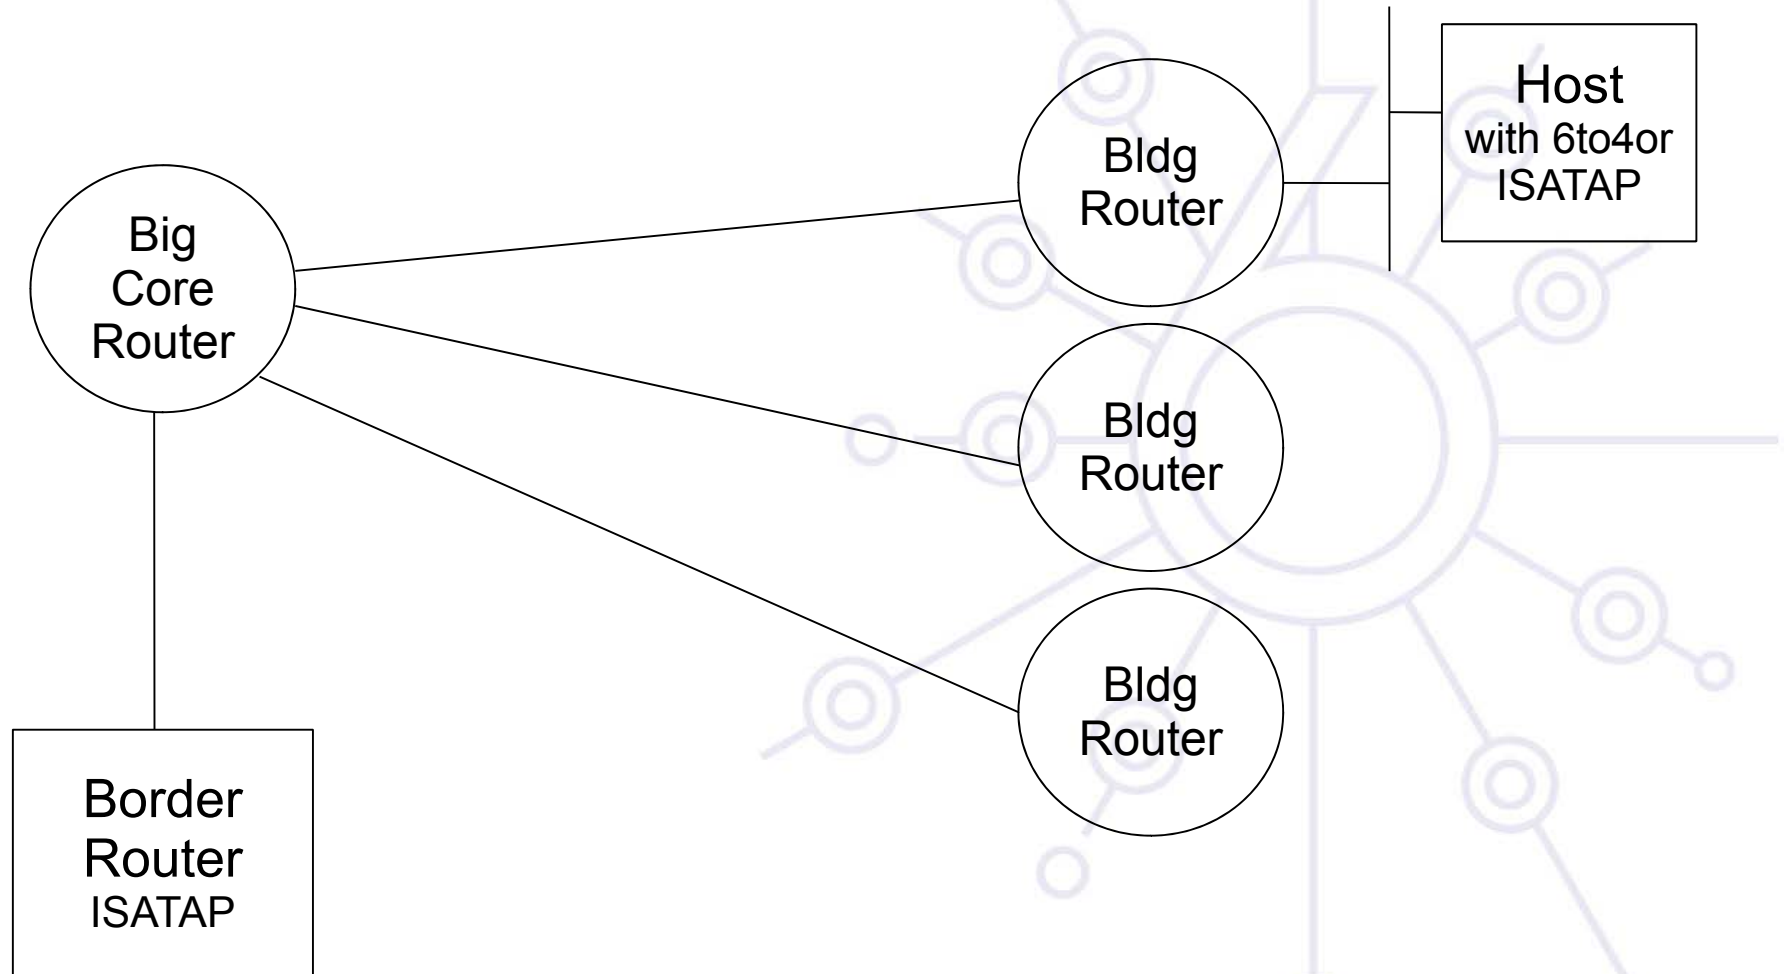
Layer-2 Campus Redundant Switches



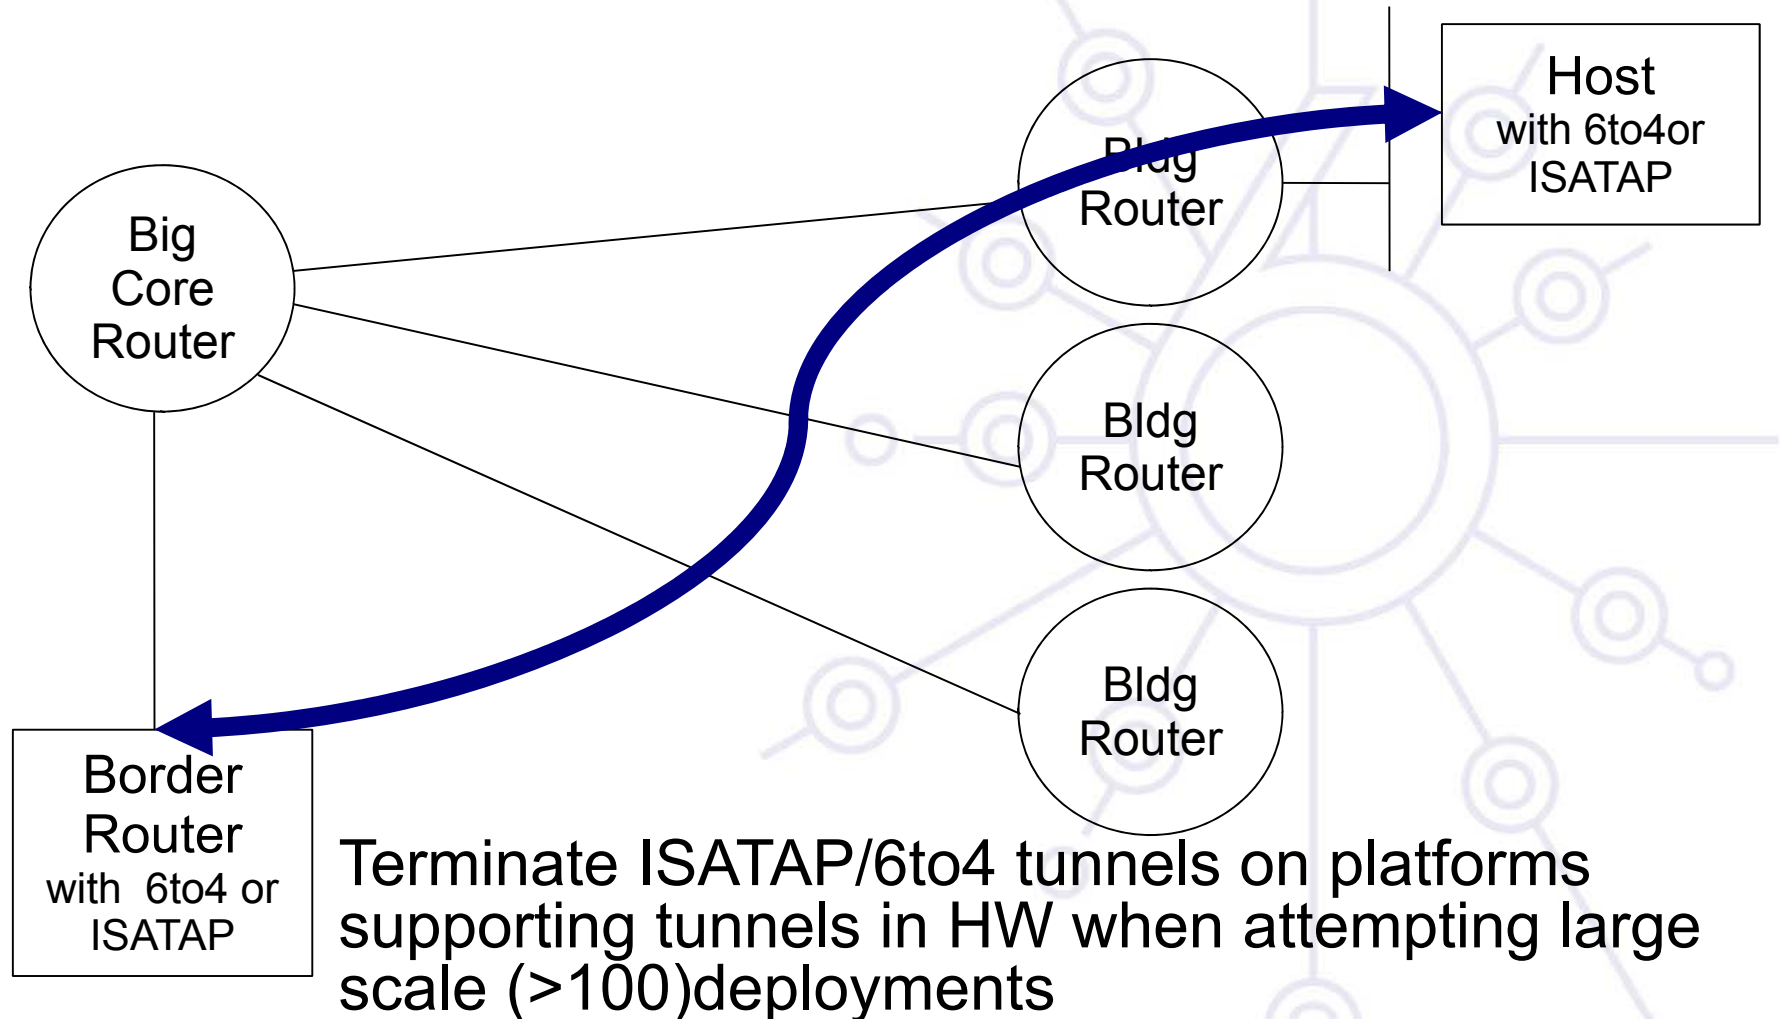
Layer-3 Campus



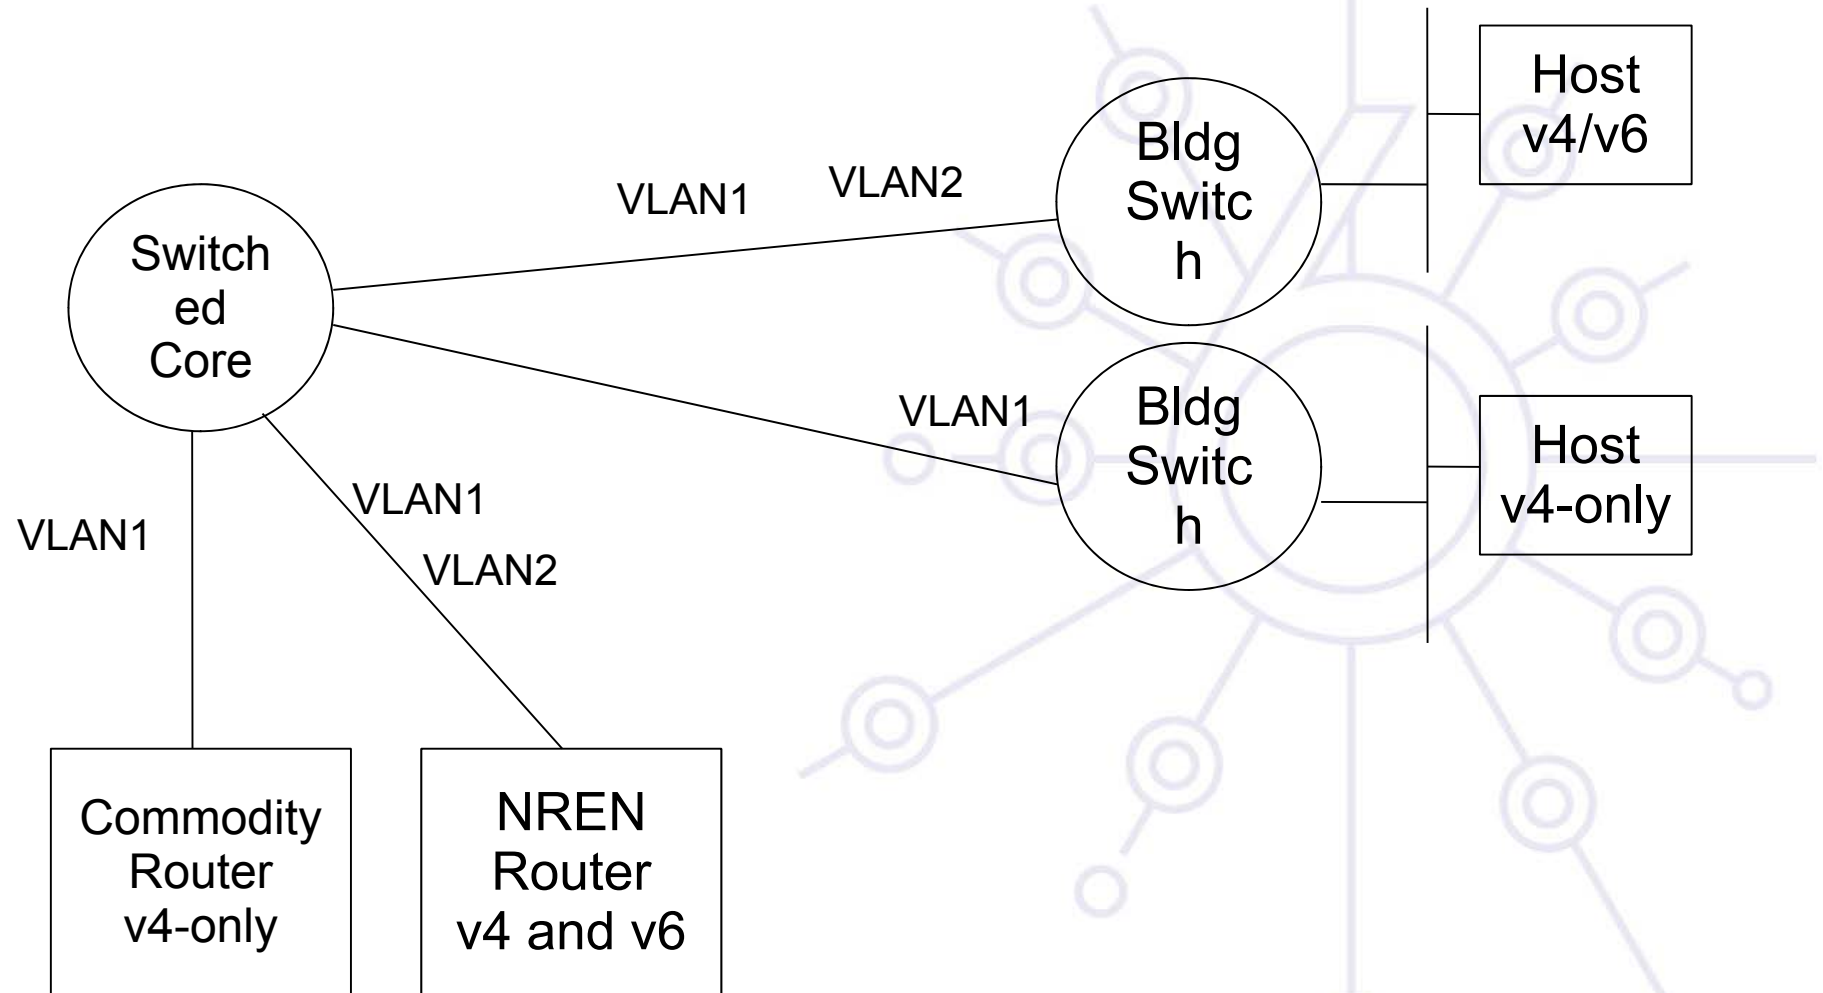
Layer-3 Campus



Layer-3 Campus



Edge Router Options



Routing Protocols

iBGP and IGP (IS-IS/OSPFv3)

- IPv6 iBGP sessions in parallel with IPv4
- You need a 32 bit router-id for IPv6 BGP peering configuration

Static Routing

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

OSPFv3 might be good

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

Use the same (type) of protocol you used in IPv4.

See more in routing module

Outline

Campus deployment strategy

Campus IPv6 address allocation

Campus deployment topology - options

Campus services

Service provider deployment considerations



Campus services –Road Map

- **Name service - see DNS module**
- **Security policy - see security module**
- **Routing - see routing module**
- **(Mail) not considered here - see application module**
- **Proxying**
- **Remote access**
- **Monitoring the network and the services - see monitoring module**

=> For most of these services, refer to the ad hoc modules on <http://www.6deploy.org>

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 an IPv6 capable server?

Use proxy (more exactly reverse-proxy) server

- Apache2.x 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

Management and monitoring

- Device configuration and monitoring -SNMP
- Statistical monitoring e.g. Cricket/MRTG
- Service monitoring - Nagios
- Intrusion detection (IDS)
- More information
 - Module #060 : IPv6 Networks management
 - <http://www.6deploy.org>

Remote access via IPv6

Use native connectivity when available

- Rather easy if you are operating dial-in pool or you are an ADSL service provider
- ... and even more easy if your home ISP provides IPv6 connectivity
 - Like Free and Nerim in France

Use (Open)VPN

Use tunnel broker service – rather suboptimal ?

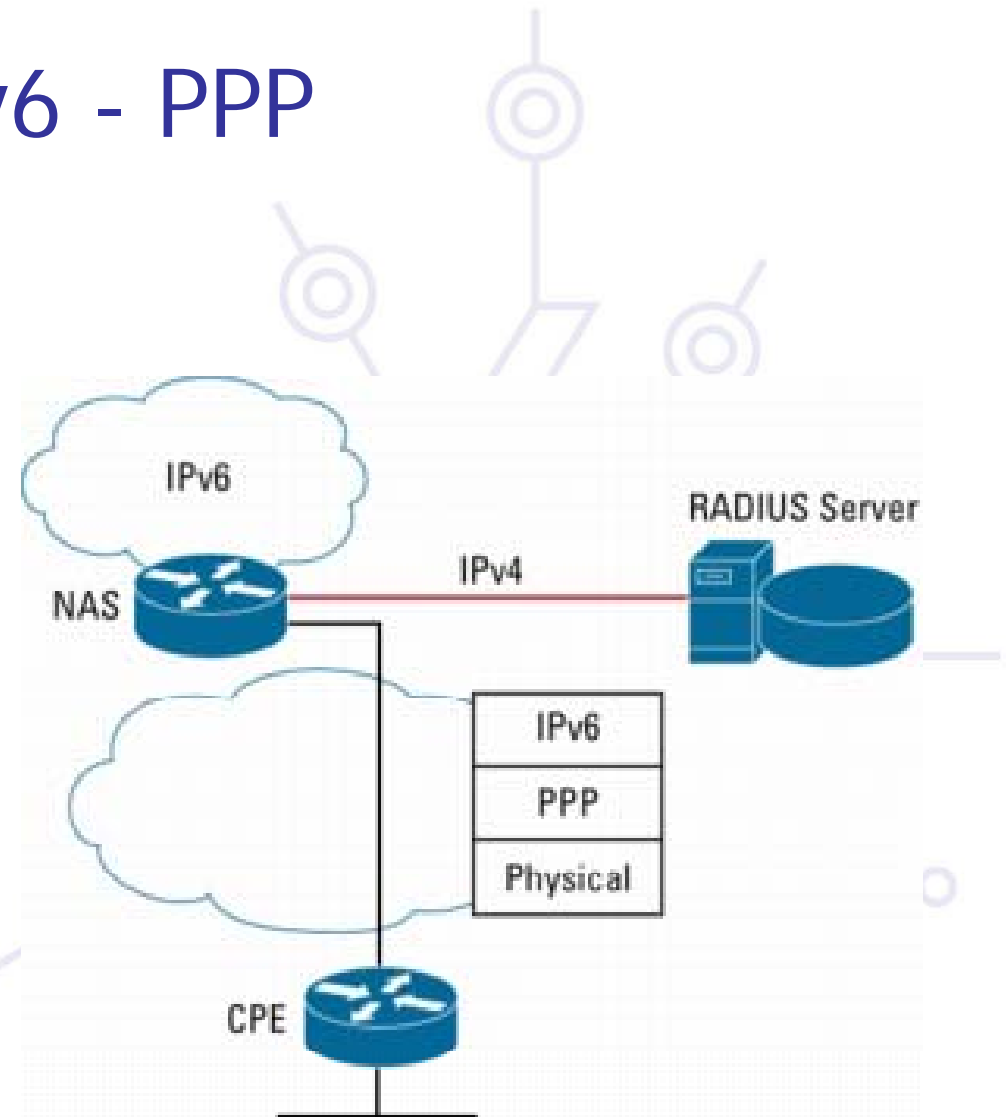
Use 6to4 if you have global IPv4 address

- Good 6to4 relay connectivity is a must

Use Teredo/software if you have NAT or multiple level of NATs

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 use PPP (Point-to-Point Protocol), which gives a standard method to transport the datagrams of several protocols over point-to-point links (RFC1661, 2153, 5342) - PPP has been updated to support the transport of IPv6 datagrams (RFC5072)



PPP and IPv6

PPP protocol has three main parts

- Definition of the encapsulation method of the IPv6 datagrams over the point- to-point link (IP6CP)
- LCP (Link Control Protocol) used to establish, configure and test the connection at link layer
- NCP (Network Control Protocol) used 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 IP6CP (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

PPP and IPv6 - implementations

Routers:

- Cisco
- Juniper

Hosts:

- Windows Vista and Microsoft Windows Server 2008
 - Windows XP: Cfos IPv6 link http://www.cfos.de/ipv6_link/ipv6_link_e.htm
- Linux, *BSD (including Mac OS X), Solaris

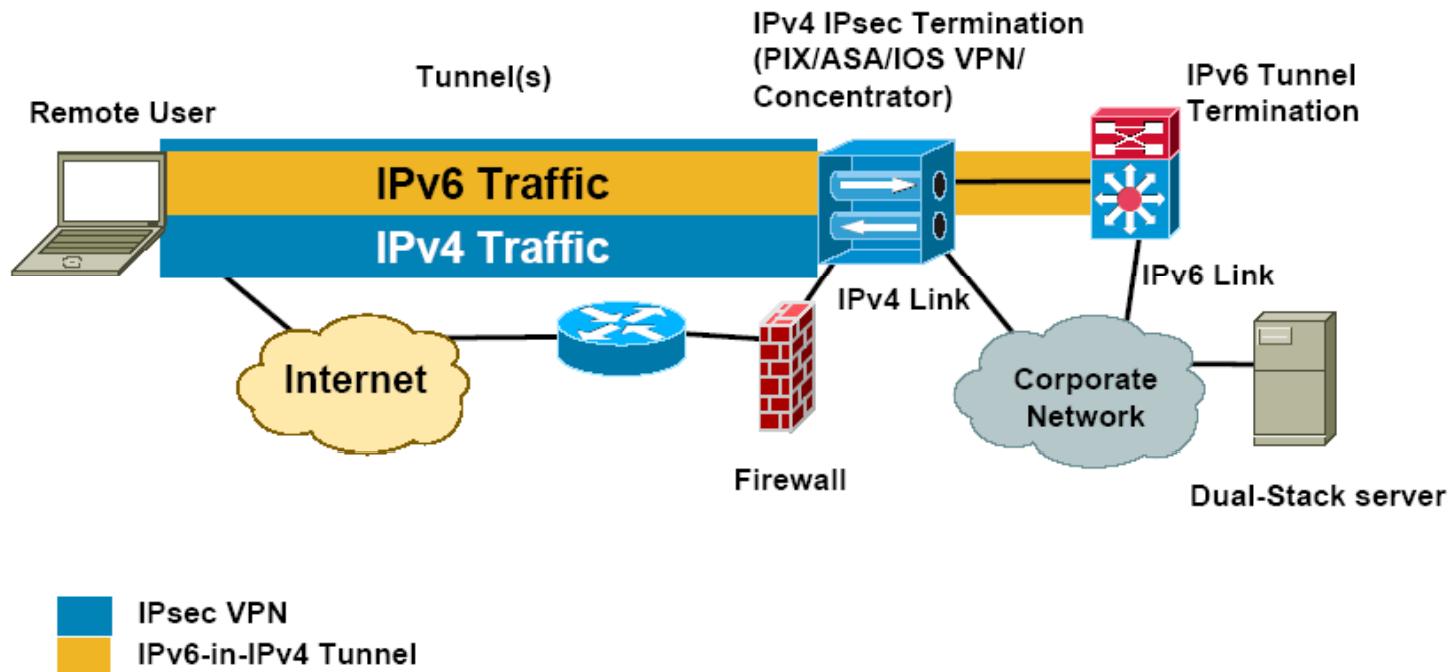
Opensource:

<http://sourceforge.net/projects/pppcbcp>

<http://freshmeat.net/projects/pppd>

Remote Access with IPSEC – or other VPNs

IPv6-in-IPv4 Tunnel Example



IPv6 load balancing

- Server clusters
 - Opensource solution: *BSD pf (<http://www.openbsd.org/faq/pf/>), Linux LVS after 2.6.28 (http://kb.linuxvirtualserver.org/wiki/IPv6_load_balancing)
 - Commercial platforms: Veritas Cluster Server, BigIron F5, Windows Server 2008 - Network Load Balancer
- First-Hop Redundancy:
 - HSRPv6 (Cisco only)
 - VRRPv6 - standardisation at IETF
 - NUD (Neighbor Unreachability detection)- 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/Datacenters ... 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



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

- Infrastructure addresses
- Routable /48 address

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
```

```
mnt-lower:     NIIF6-MNT
```

```
status:        ALLOCATED-BY-RIR
```

← New mandatory
← New mandatory
← New

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 is in production (at least in RIPE region)

Reverse delegation is strongly recommended

Summary

Campus deployment strategy

- Coexistence mechanism ?
- Getting an IPv6 prefix
- ... and external IPv6 connectivity
- Decide a security policy for IPv6 traffic

Campus IPv6 address allocation and usage

- Work out an addressing plan
- Decide which address allocation mechanism will be used

Campus deployment topology - options

- Start IPv6 deployment
- How to remote access the campus ?

Campus services

- Enable services for IPv6
 - Starting with the DNS
- Enable management and monitoring tools
- Enable IPv6 on hosts



deploy

Questions ...