IPv6 transition – IPv6 deployment

János Mohácsi

NIIF/HUNGARNET

Central Asia Workshop

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*

Contributions

- Main authors
  - János Mohácsi, NIIF/HUNGARNET - Hungary
- Contributors
  - Jérôme Durand, Renater, France
  - Tim Chown, University of Southampton, Great-Brittain

Various Campus transition approaches

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

Tunnelling

- Initially IPv6 in IPv4 or IPv6 in MPLS, (much) later IPv4 in IPv6
- So, IPv6 packets are encapsulated in IPv4 packets/MPLS frame
  - IPv6 packet is payload of IPv4 packet/MPLS frame
- Usually used between edge routers to connect IPv6 'islands'
  - Edge router talks IPv6 to internal systems
  - Encapsulates IPv6 in IPv4/MPLS towards remote tunnel endpoint

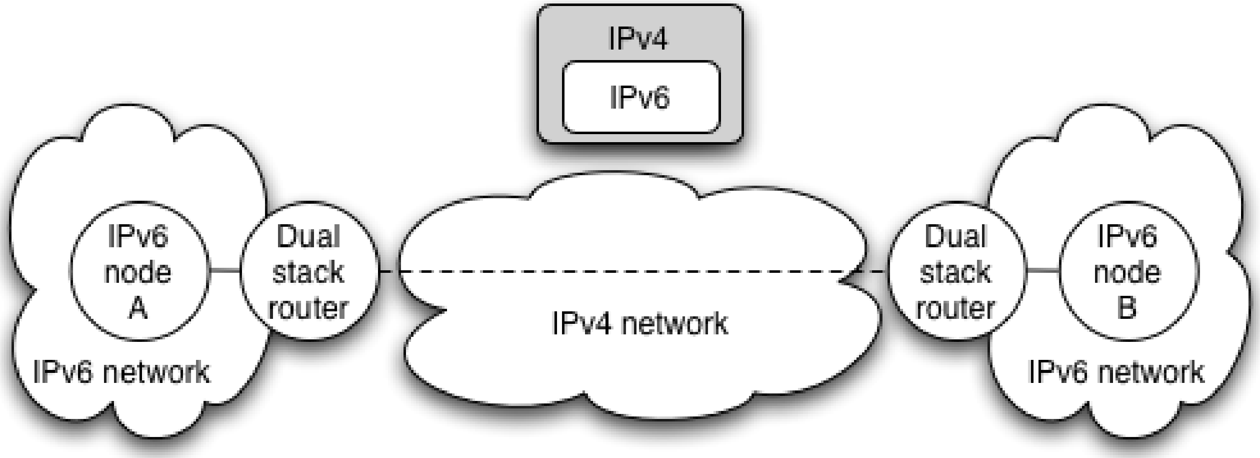

6to4

- In its basic configuration, 6to4 is used to connect two IPv6 islands across an IPv4 network
- Uses special ‘trick’ for the 2002::/16 IPv6 prefix that is reserved for 6to4 use
  - Next 32 bits of the prefix are the 32 bits of the IPv4 address of the 6to4 router
  - For example, a 6to4 router on 192.0.1.1 would use an IPv6 prefix of 2002:c000:0101::/48 for its site network
- When a 6to4 router sees a packet with destination prefix 2002::/16, it knows to tunnel the packet in IPv4 towards the IPv4 address indicated in the next 32 bits

6to4 basic overview

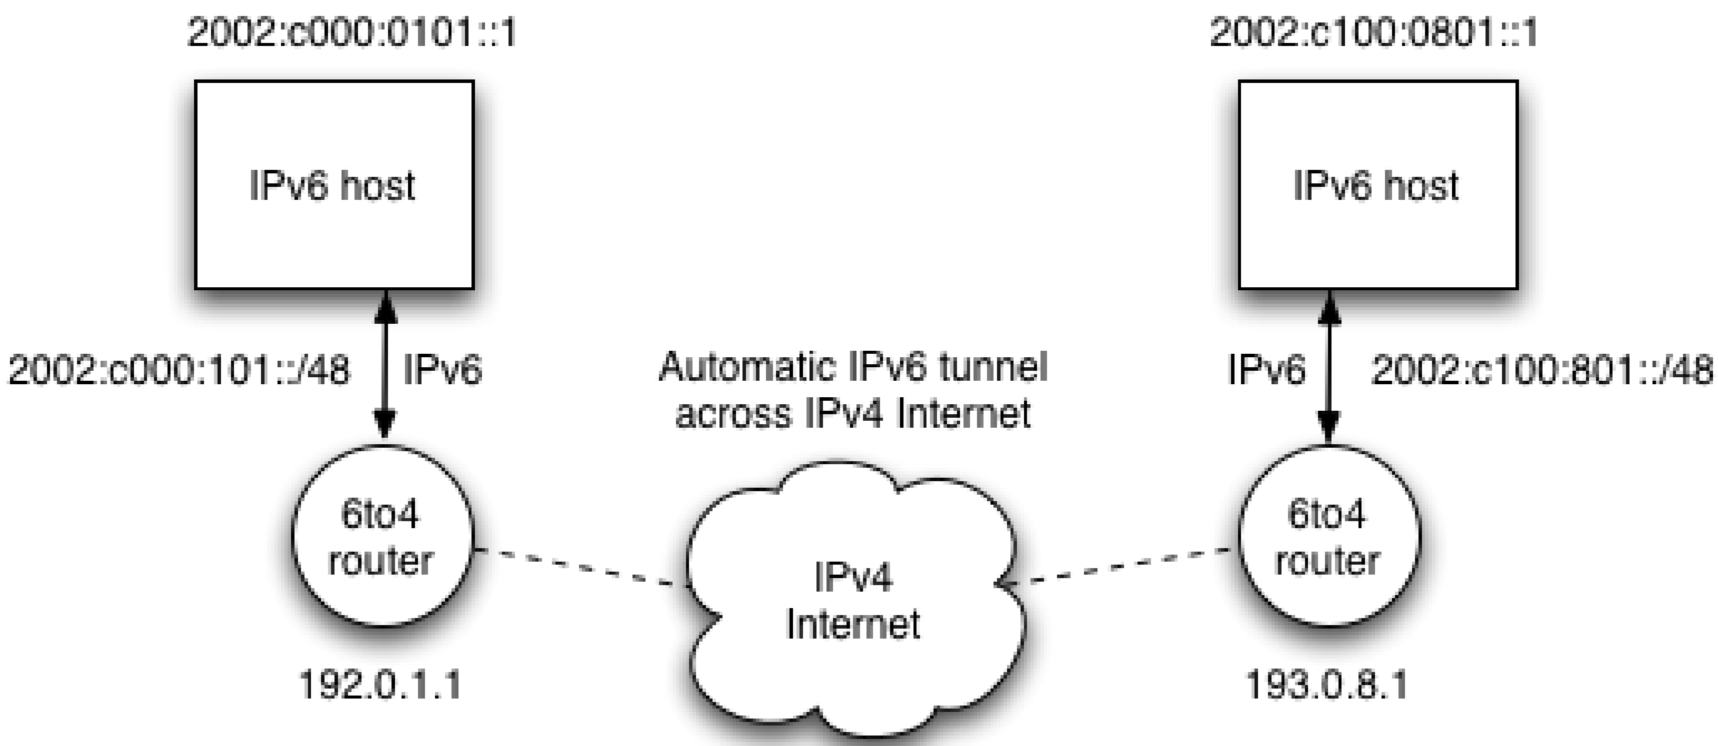

6to4 with relay

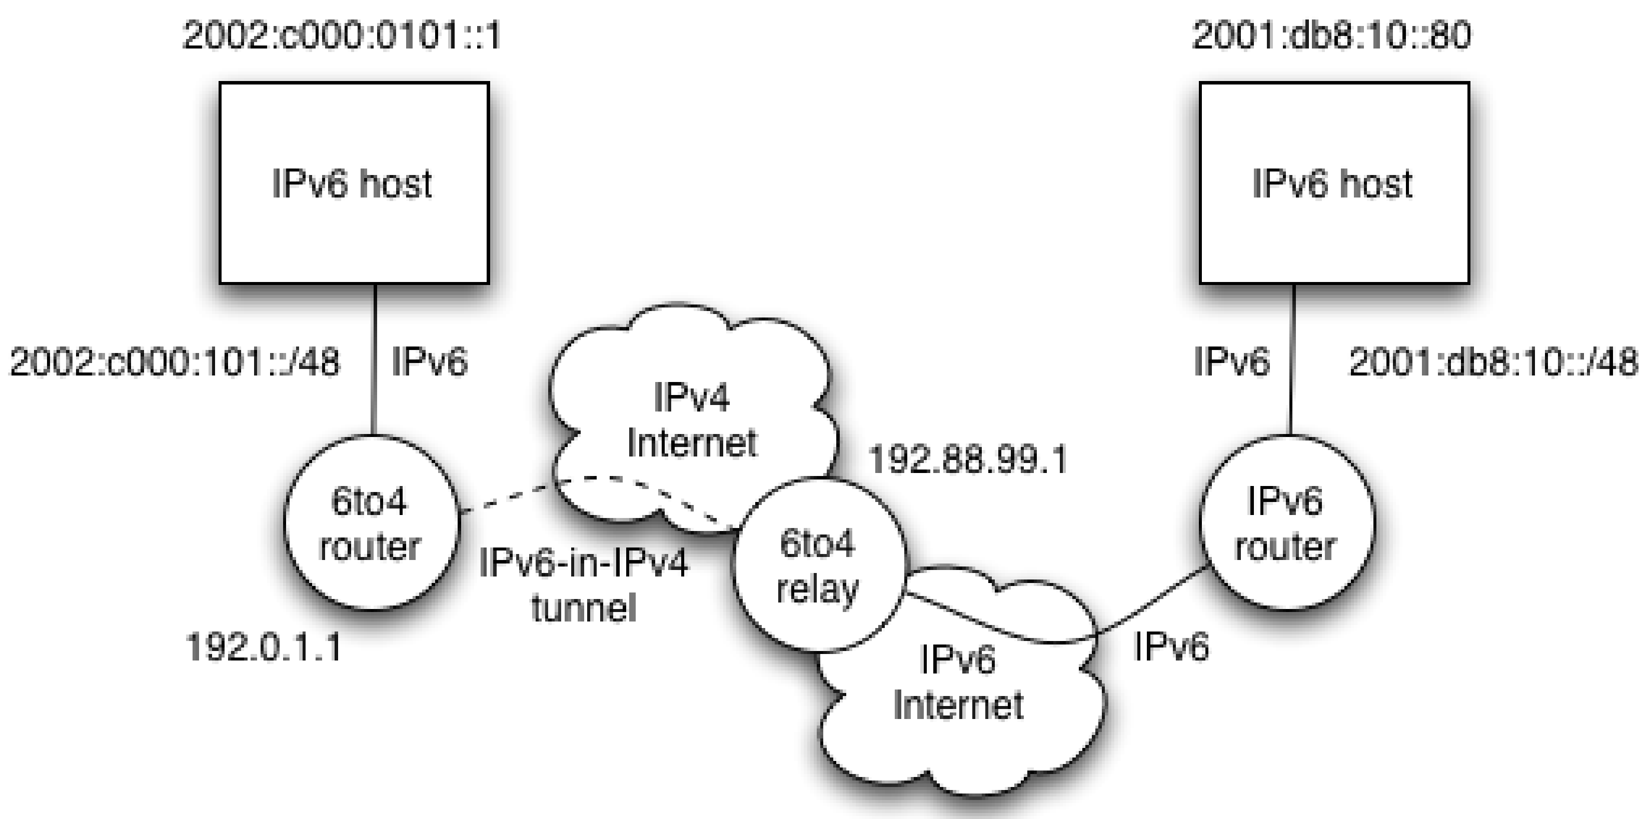

ISATAP

- Intra-Site Automatic Tunnel Addressing Protocol (RFC4214)
  - Automatic tunneling
  - Designed for use *within* a site
  - Used where dual-stack nodes are sparsely deployed in the site (very early deployment phase)
- Host-to-host or host-to-router automatic tunnels
  - Uses a specific EUI-64 host address format
  - Format can be recognised and acted upon by ISATAP-aware nodes and routers

ISATAP addresses

- The EUI-64 is formed by
  - A reserved IANA prefix (00-00-5e)
  - A fixed 8-bit hex value (fe)
  - The 42-bit IPv4 address of the node
  - Toggling the globally unique (u) bit
- For example, 152.66.64.1 would have an EUI-64 host address for IPv6 of:
  - 0200:5efe:9842:4001

ISATAP tunneling

- Relies on the OS supporting ISATAP
- Use one ISATAP router per site, usually advertised under FQDN ‘isatap.domain’
  - Virtual IPv6 link over the IPv4 network
  - Know the IPv4 tunnel end-point address from last 32-bits of the IPv6 ISATAP address
  - Get network prefix via ND from router
- Not widely deployed
- Better to deploy proper dual-stack
  - Allows better managed control of deployment

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

Outline of NRENs/ISP IPv6 deployment

- Obtain IPv6 address space
- Plan the addressing
- Plan the routing
- Test in a small case
- Deploy IPv6 (incrementally – dual-stack/6PE)
- 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

Addressing architecture at NIIF/HUNGARNET (case study)

János Mohácsi
janos.mohacsi(a)niif.hu

Site addressing

- Each site (including site infrastructure) get /48 but future extensibility to /44:
  - each NIIF managed site the 16 bit SLA is allocated based on the following convention: <SLA> = Address segmentation within the POP
  - Where for <SLA>:
    - Range: 0000 till 00FF: Loopback addresses
    - Range: 0100 till 01FF: Intra-pop point-to-points (if it necessary to number it)
    - Range: 0200 till 02FF: connections to HUNGARNET member of institution
    - Range: 0300 till 03FF: external connectivity (e.g. peering)
    - Range: 0400 till 04FF: POP Local Ethernets

IPv6 loopback addresses

- loopback address will also be used for operational and management actions on the equipment, and for routing protocols like iBGP, which will use these addresses for terminating the peering-sessions.
- Loopback addresses have typically a prefix mask of /128. This will avoid unnecessary unused addresses although address conservation is not really an issue in IPv6.

p2p Link addresses?

- Not necessary!
  - OSPFv3 is working with link-local
  - For IS-IS not necessary working with CLNS
- IGP table can quite small! – helps on convergence!
- Customer network is propagated into BGP (even if static routes are used
  - but not with redistribute
  - but network)

Campus deployment plan /1

- Planning
  - IPv6 capability is a requirement
  - IPv6 training

Campus deployment plan /2

- Obtain global IPv6 address space from your ISP
  - NRENs usually has a /32 prefix from RIPE NCC/RIRs
  - A university will get a /48 prefix from NRENs
- 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

- 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 on host systems (Linux, WinXP, …)
  - Enable IPv6 services and applications
    - Starting with DNS
  - Advertise IPv6 services
  - Enable management and monitoring tools

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?

Campus Addressing

- Sequentially, e.g.
  - 0000
  - 0001
  - …
  - FFFF

  - 16 bits = 65535 subnets

Campus Addressing

- 2. Following existing IPv4:
  - Subnets or combinations of nets & subnets, or VLANs, etc., e.g.
  - `152.66.`**`60`**`.0/24` .003c
  - `152.66.`**`91`**`.0/24` .005b
  - `152.66.`**`156`**`.0/24` .009c

Campus Addressing

- 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 subnets” (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!

New Things to Think About

- 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

- 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

Interface-ID Selection – some thoughts

- Scanning - the search for something to attack
- Use random 64-bit interface-IDs
  - 2001:DB8:BEEF:2::1/64 - Common IID
  - 2001:DB8:BEEF:2::9A43:BC5D/64 - Random IID
  - 2001:DB8:BEEF:2::A001:1010/64 - Semi-random IID
- Operationally can be difficult this type of numbering scheme

Topology Issues

V6 in a production network

Dual stack

- Obviously the prefered methods
- Requires switching/routing platforms to support hardware based forwarding for IPv4 and IPv6
- IPv6 is transparent on L2 switches except for multicast - MLD snooping
- IPv6 management
  – Telnet/SSH/HTTP/SNMP
- Requires robust control plane for both IPv4 and IPv6
- Requires support for Ipv6 multicast, QoS, infrastructure security, etc...

Layer-2 Campus - one switch

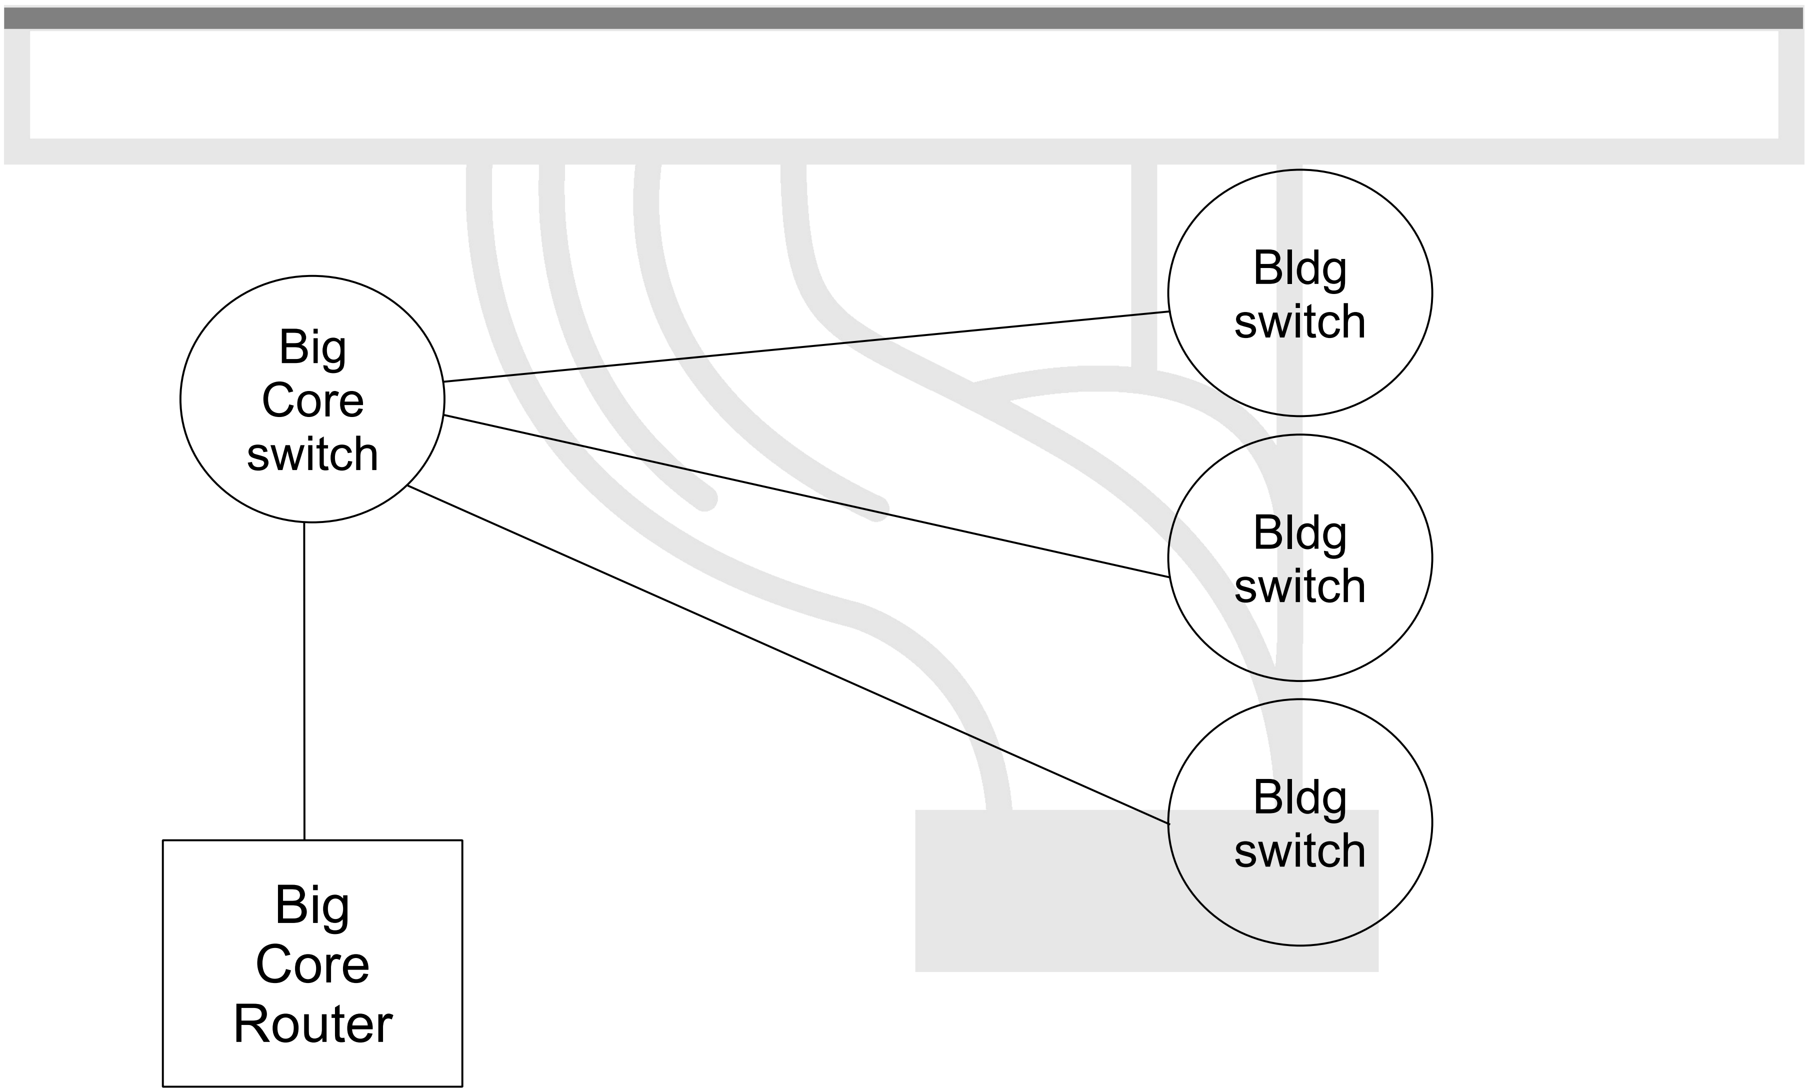

Layer-2 Campus - one Switch

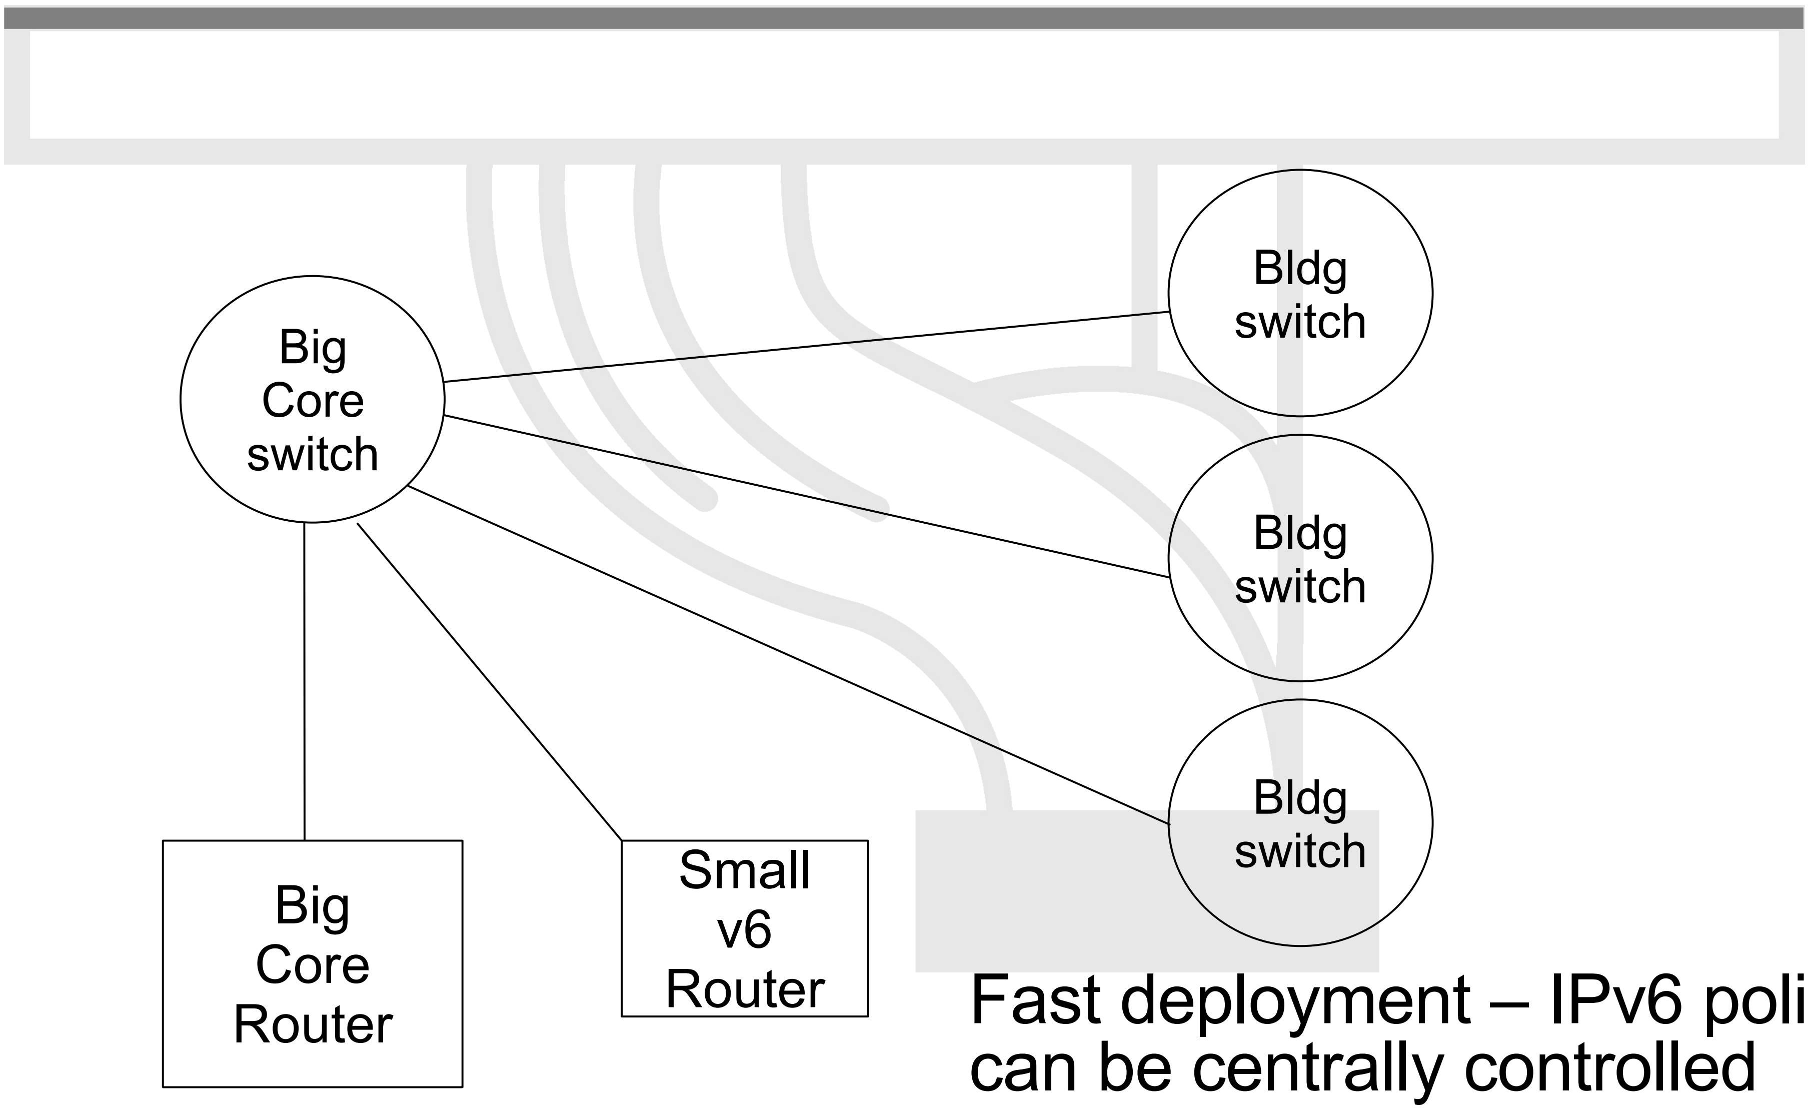

Fast deployment – IPv6 policy can be centrally controlled

Layer-2 Campus - Redundant Switches

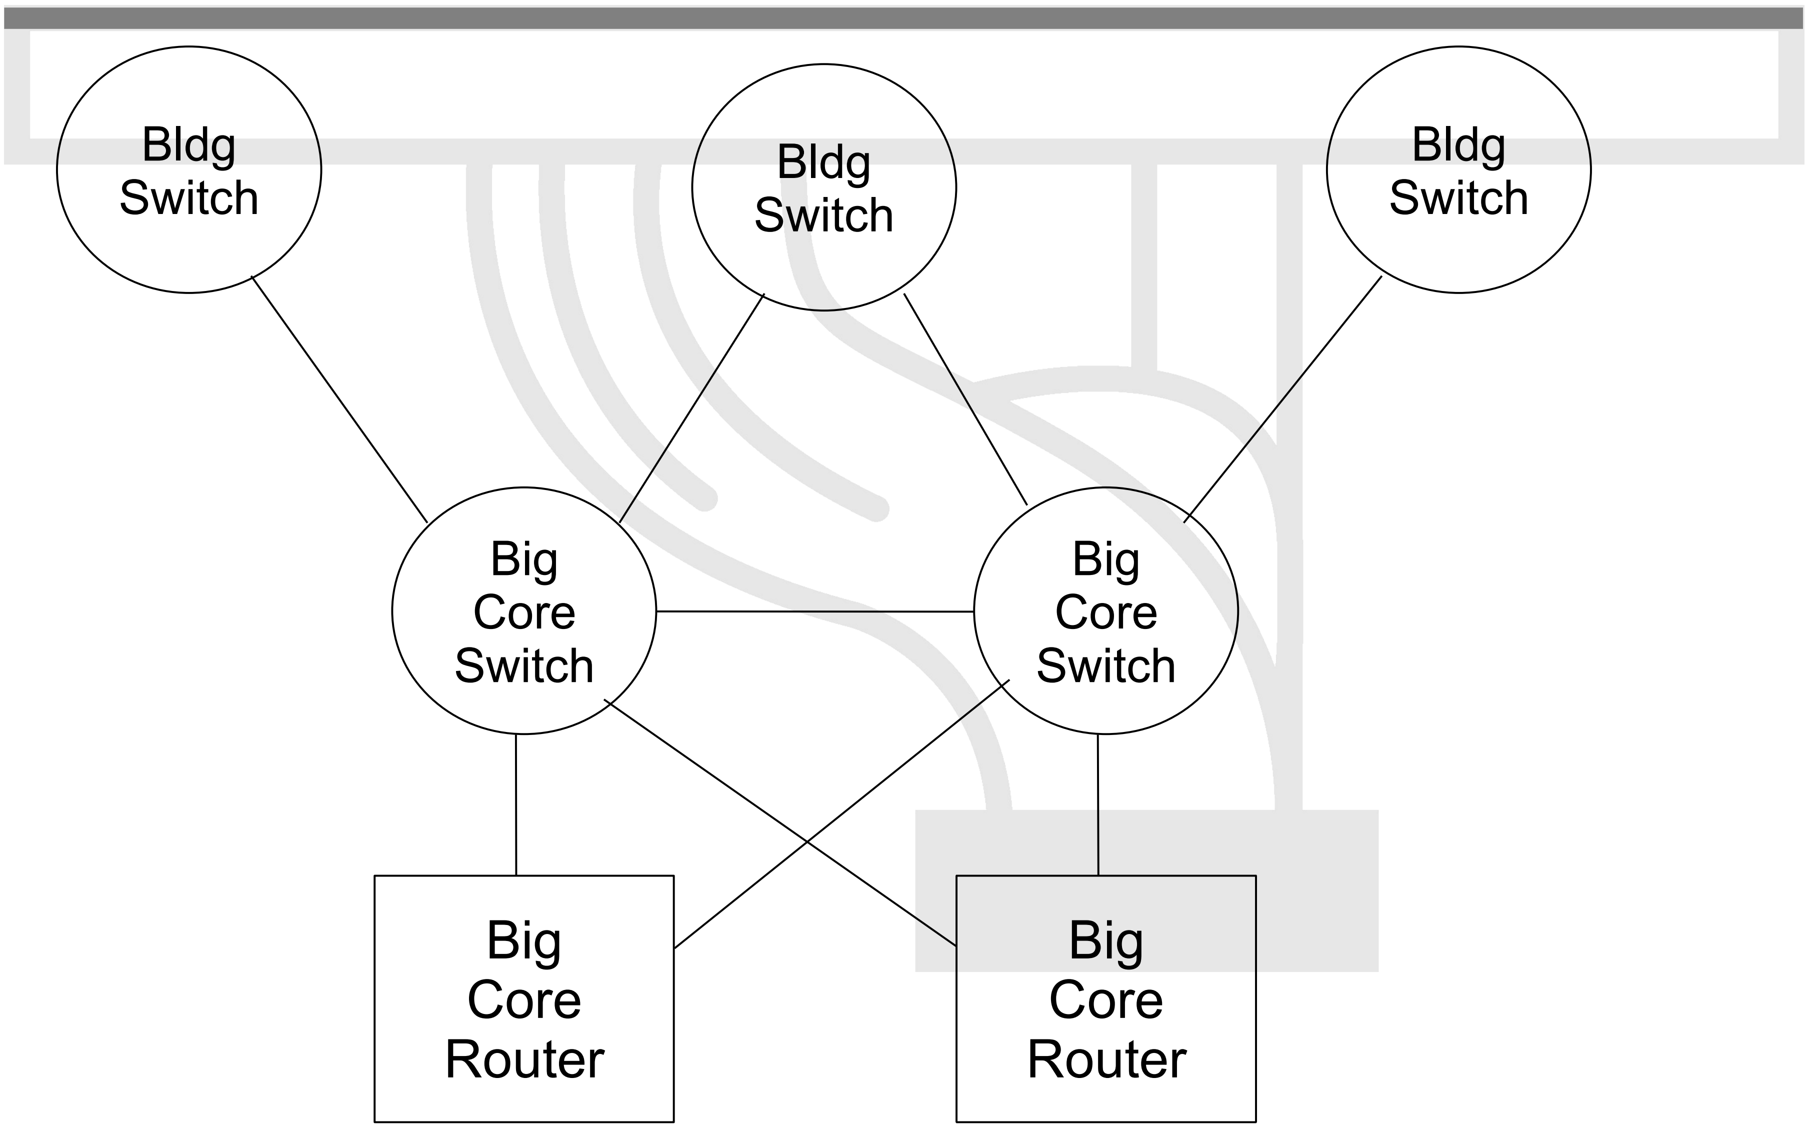

Layer-2 Campus Redundant Switches

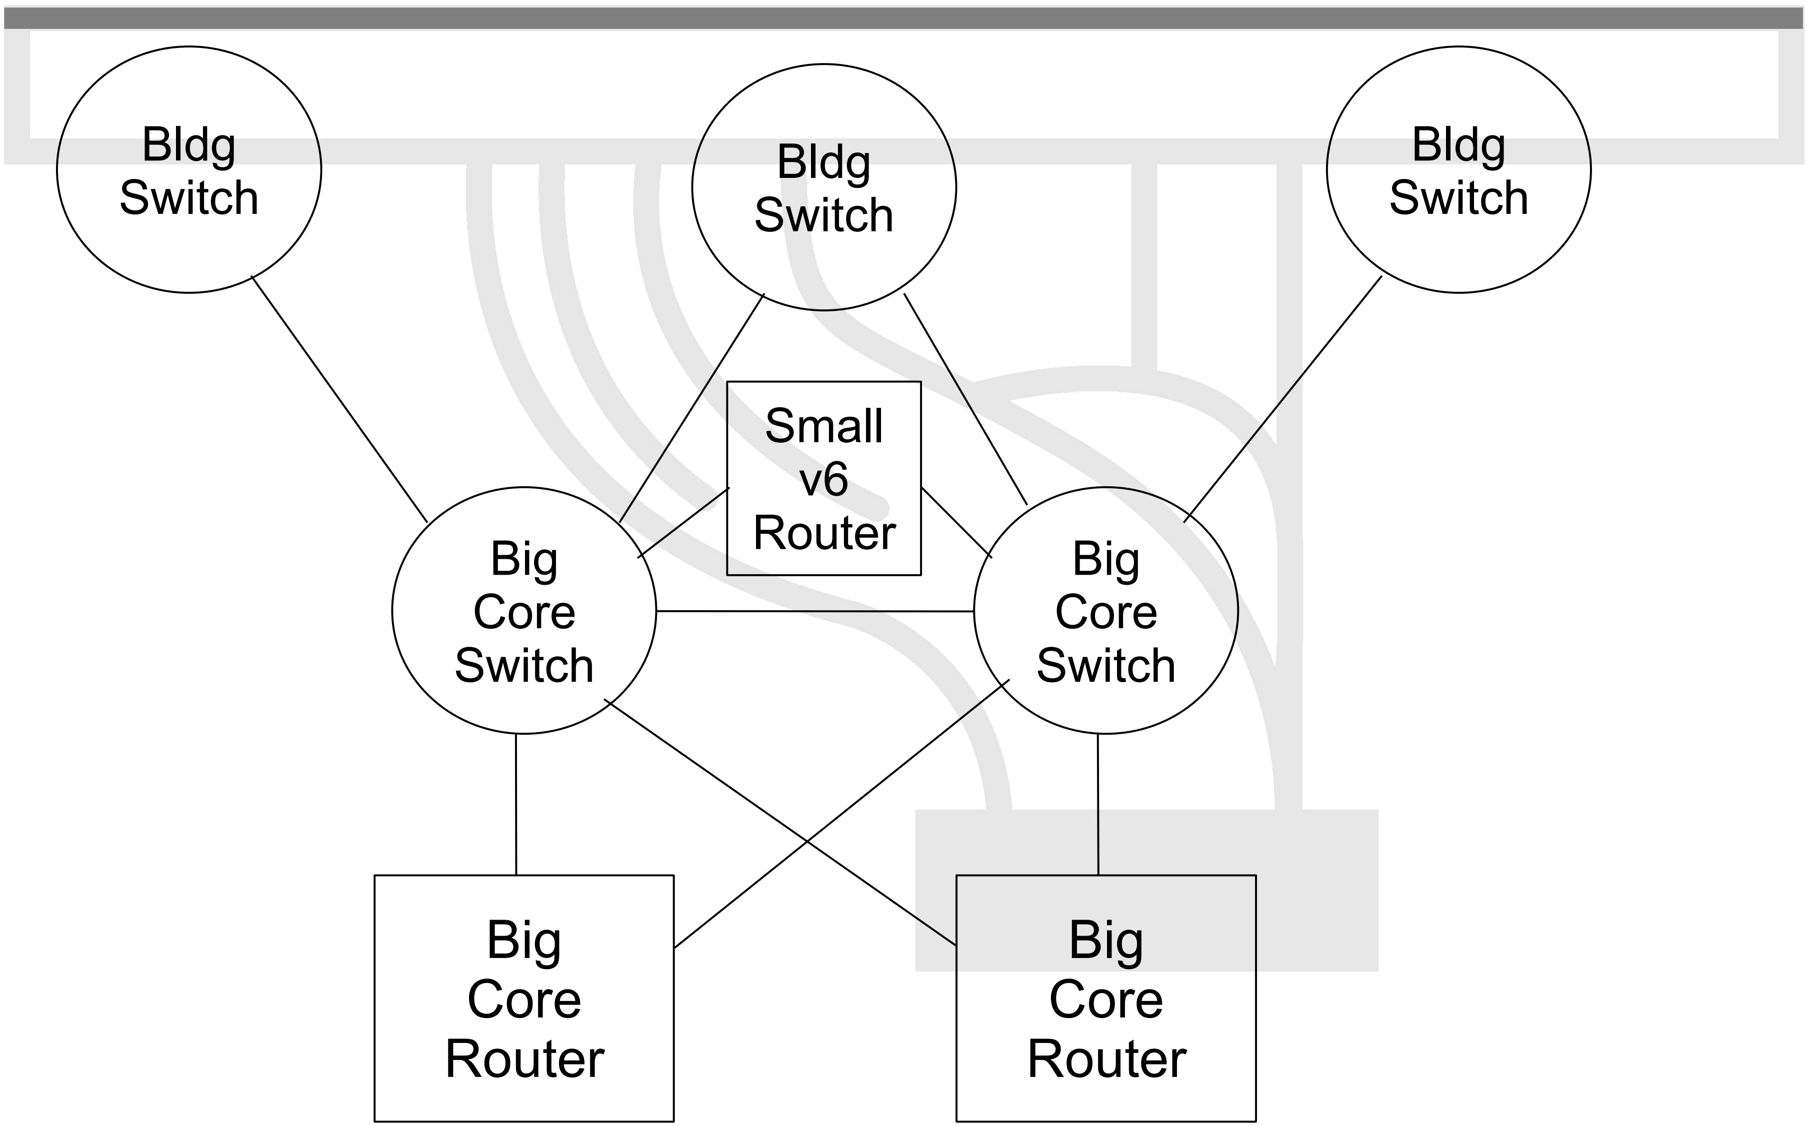

Layer-3 Campus

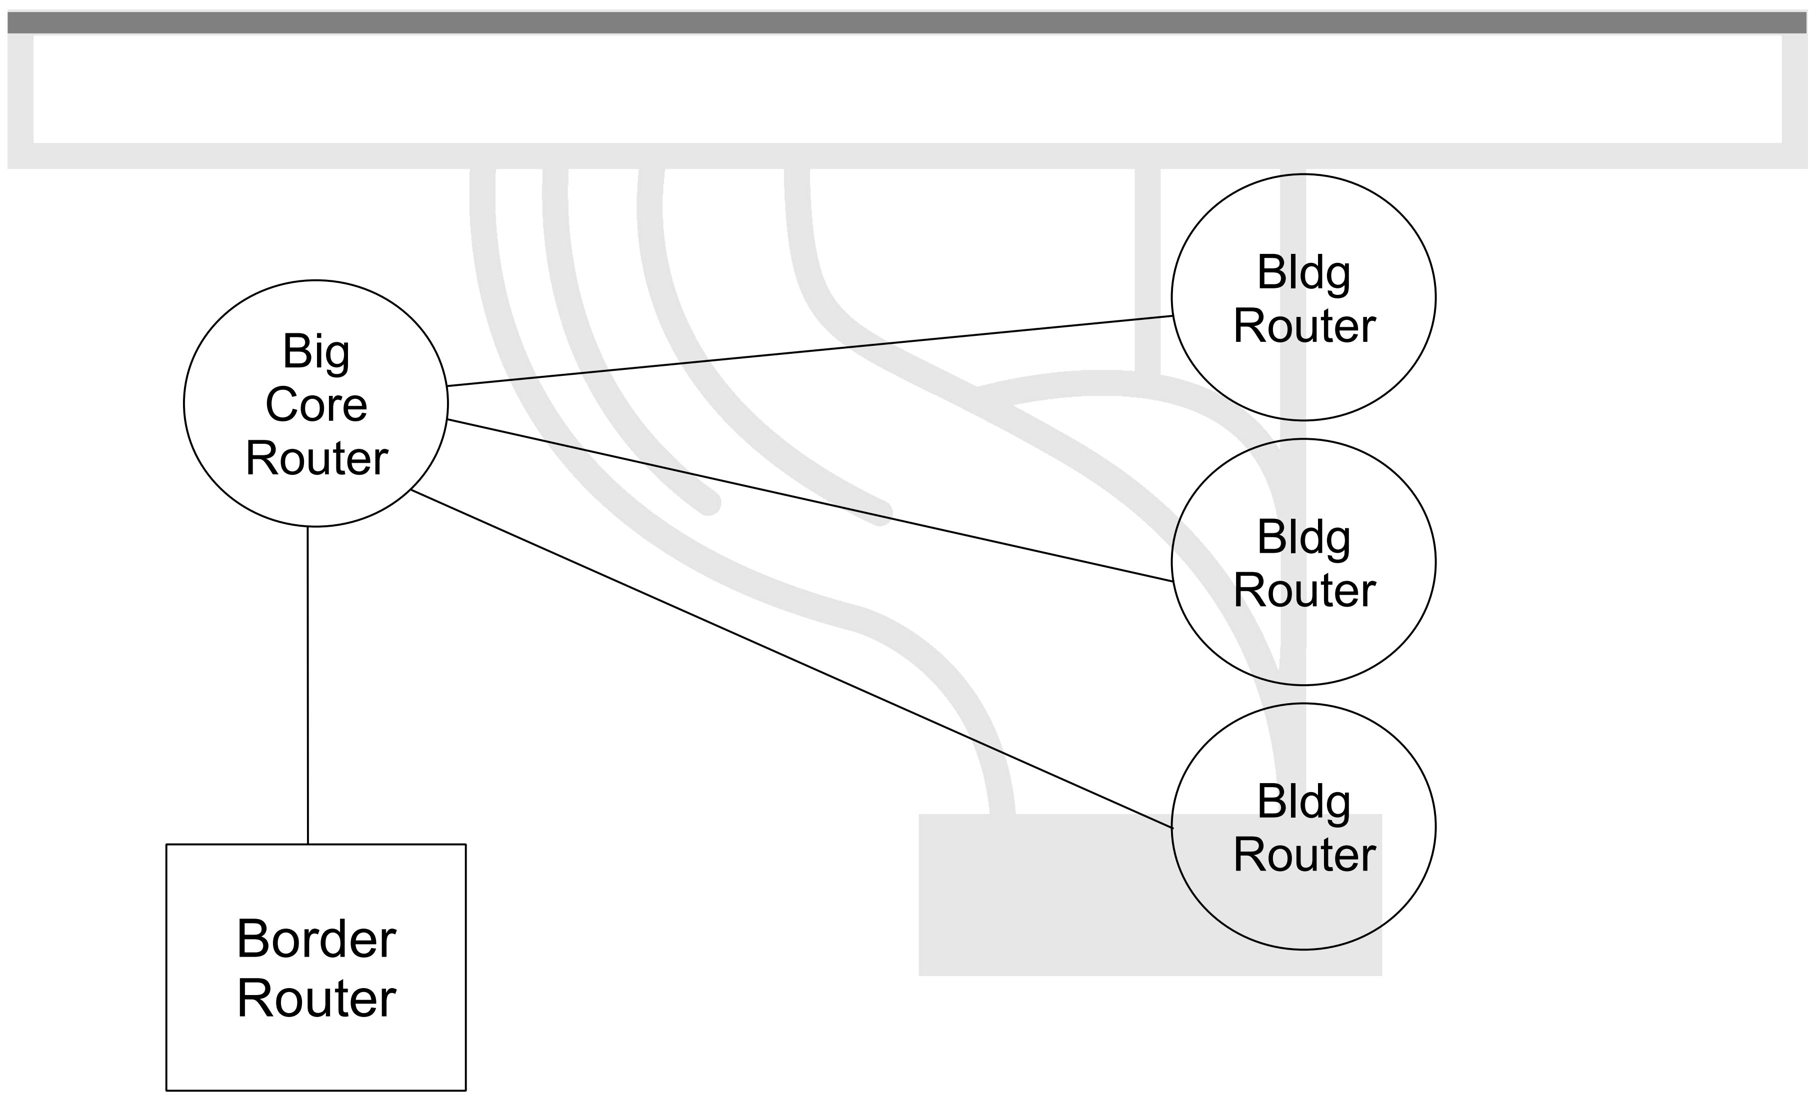

Layer-3 Campus

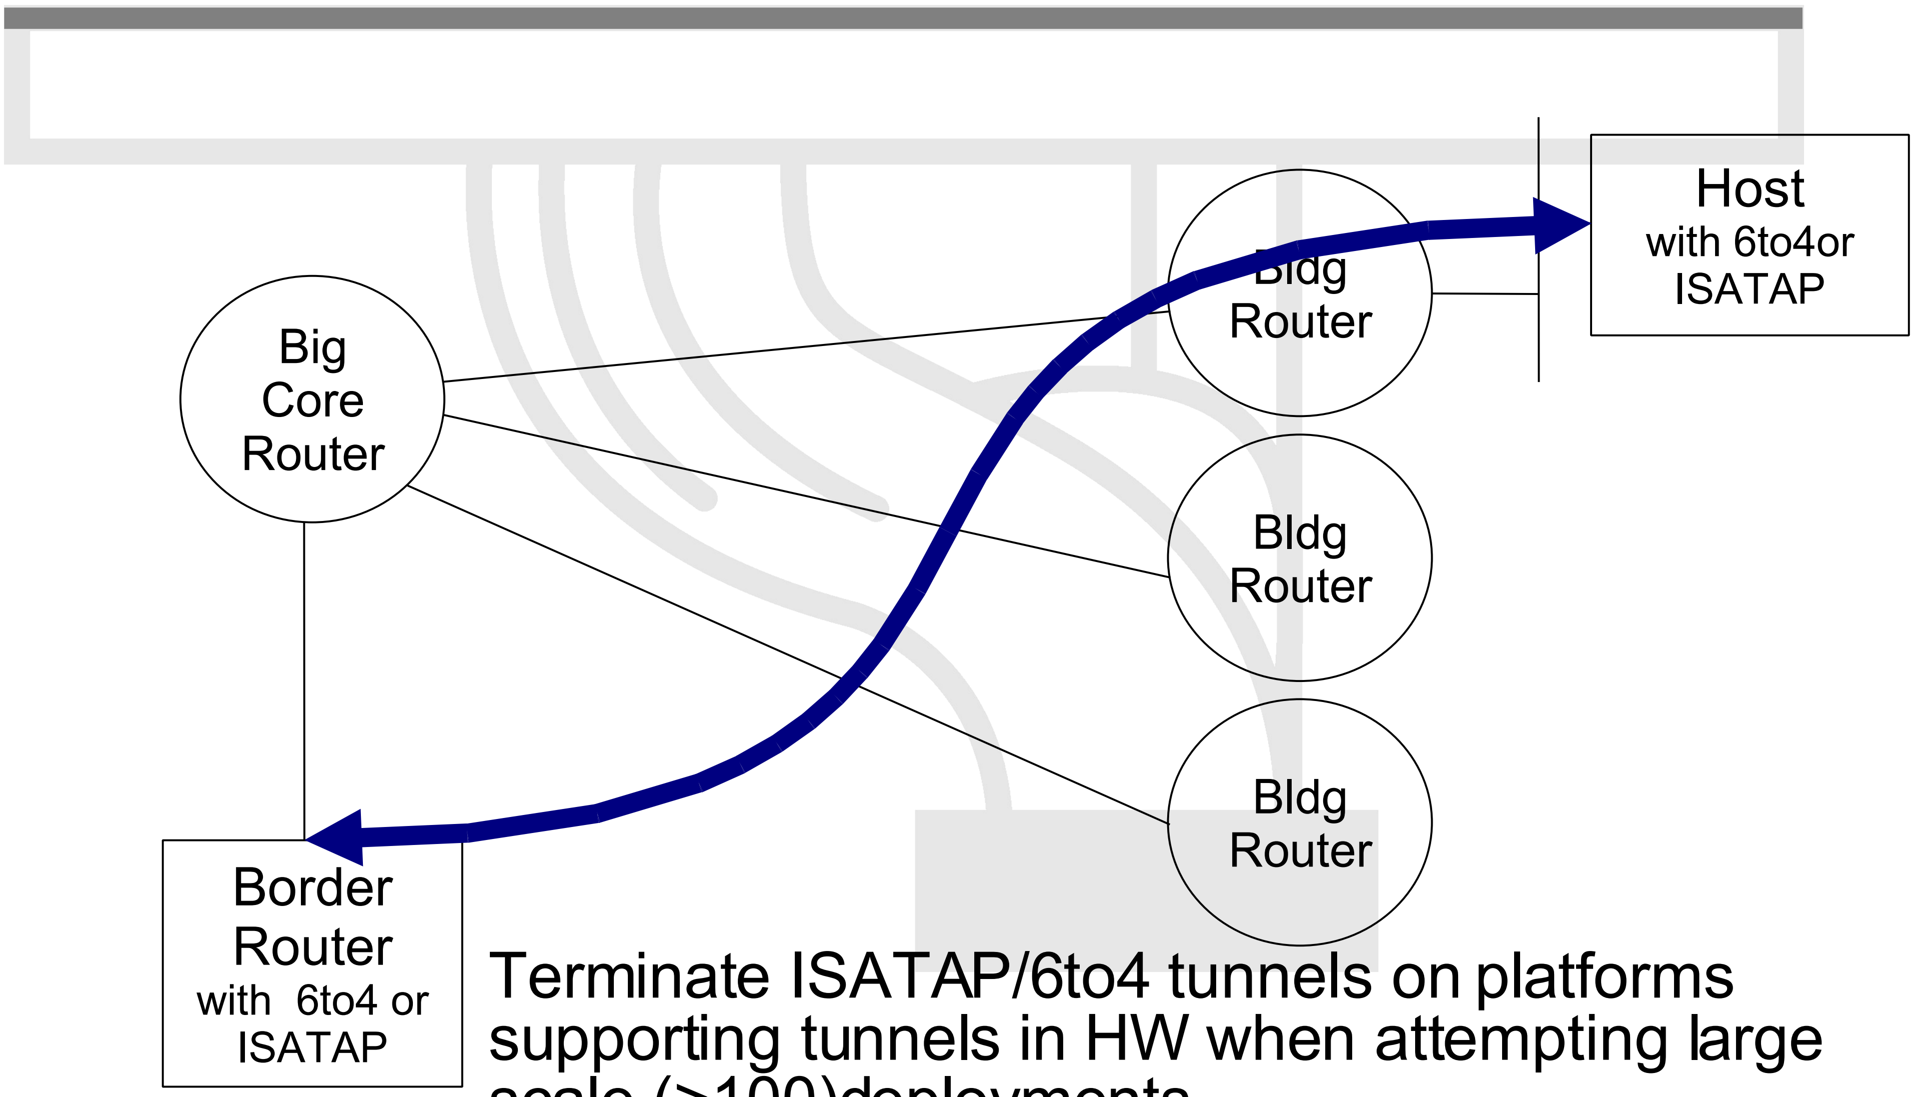

Terminate ISATAP/6to4 tunnels on platforms supporting tunnels in HW when attempting large scale (>100)deployments

Edge Router Options

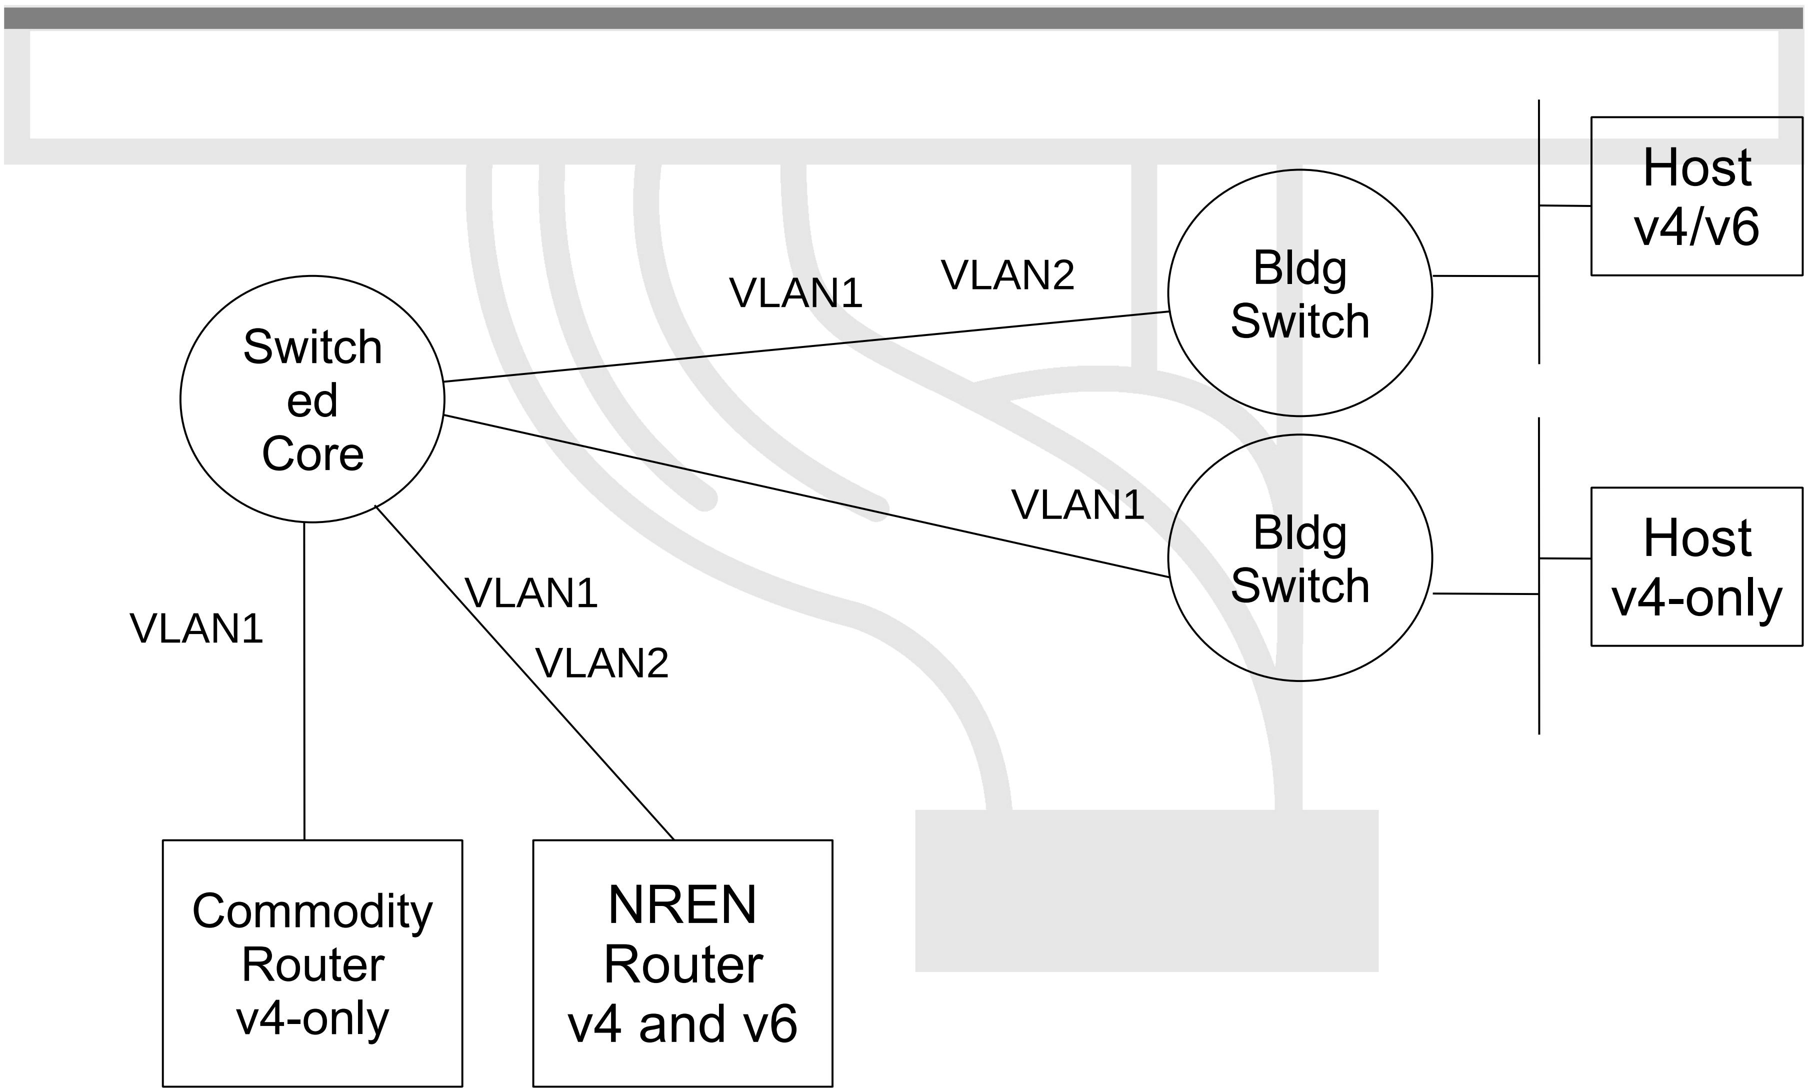

Routing Protocols

- 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.

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

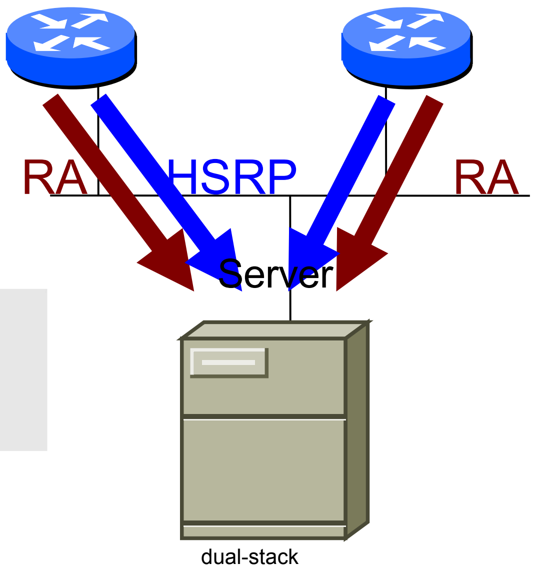

Management and monitoring

- Device configuration and monitoring -SNMP
- Statistical monitoring e.g. Cricket/MRTG
- Service monitoring - Nagios
- Intrusion detection (IDS)
- Authentication systems
  - For example, 802.1x + RADIUS for WLAN
- 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 ☺

Apache2 reverse proxy

- Configuration is very easy:

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

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

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

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

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 OpenVPN

Remote Access with IPSEC – or other VPNs

IPv6-in-IPv4 Tunnel Example

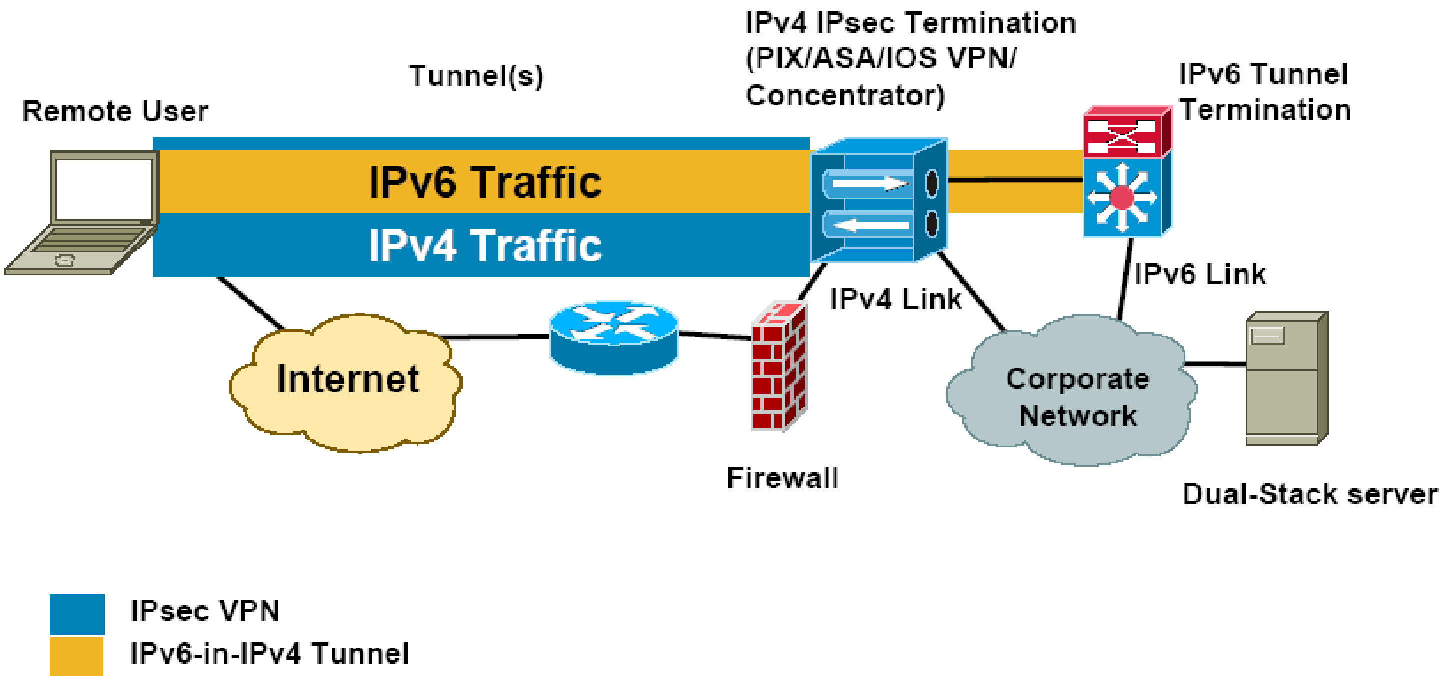

Questions?

mohacsi@niif.hu