Application Deployment considerations

János Mohácsi
NIIF/HUNGARNET

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
  - Miguel Baptista, FCCN, Portugal
  - Carlos Friaças, FCCN, Portugal
  - Laurent Toutain, ENST-Bretagne – IRISA, France
  - Bernard Tuy, Renater, France
  - Jérôme Durand, Renater, France
  - Tim Chown, University of Southampton, UK

Agenda

- Deploying IPv6 campus networks
  - Strategies,Topology, addressing,
- Basic IPv6 network services
  - DNS, other basic network applications

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

Campus deployment plan /1

1. 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
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

1. 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

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

Topology Issues

V6 in a production network

Layer-2 Campus -1 Switch

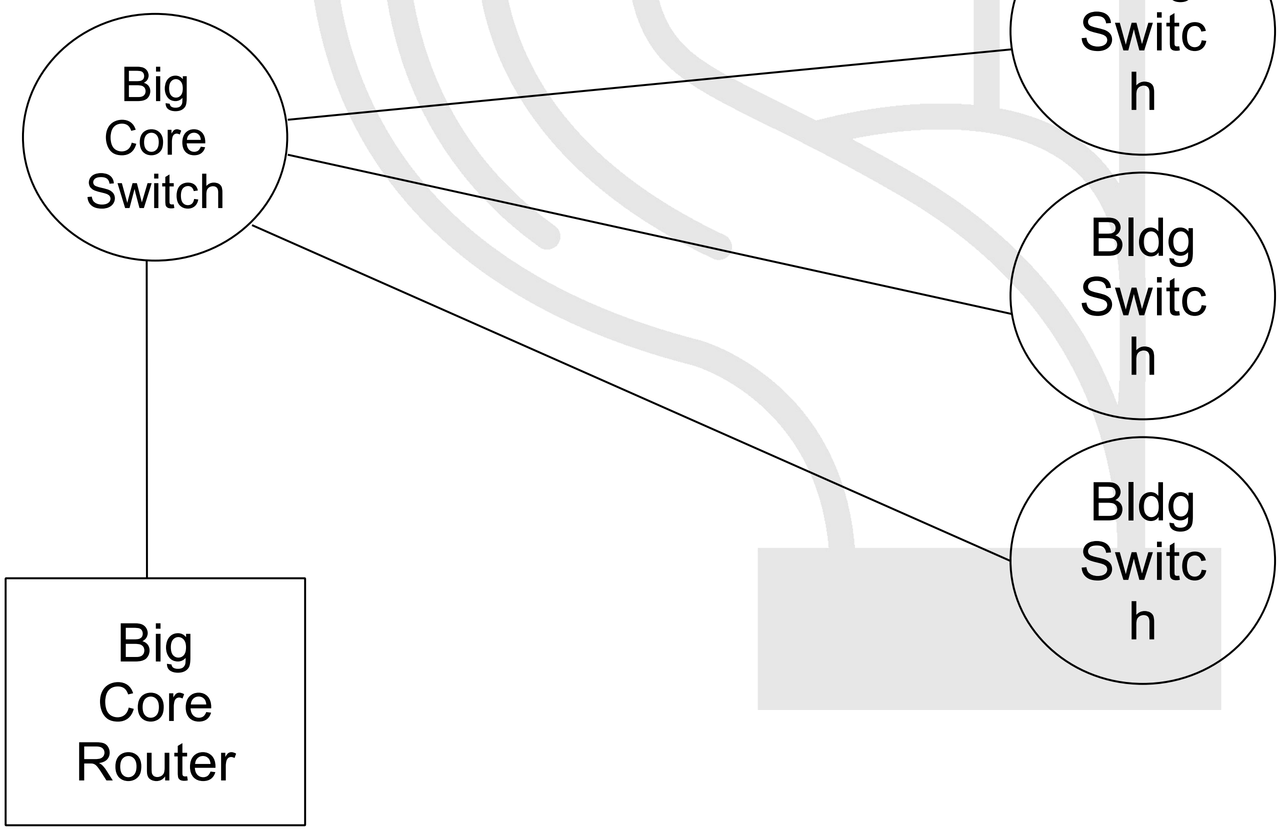

Layer-2 Campus - 1 Switch

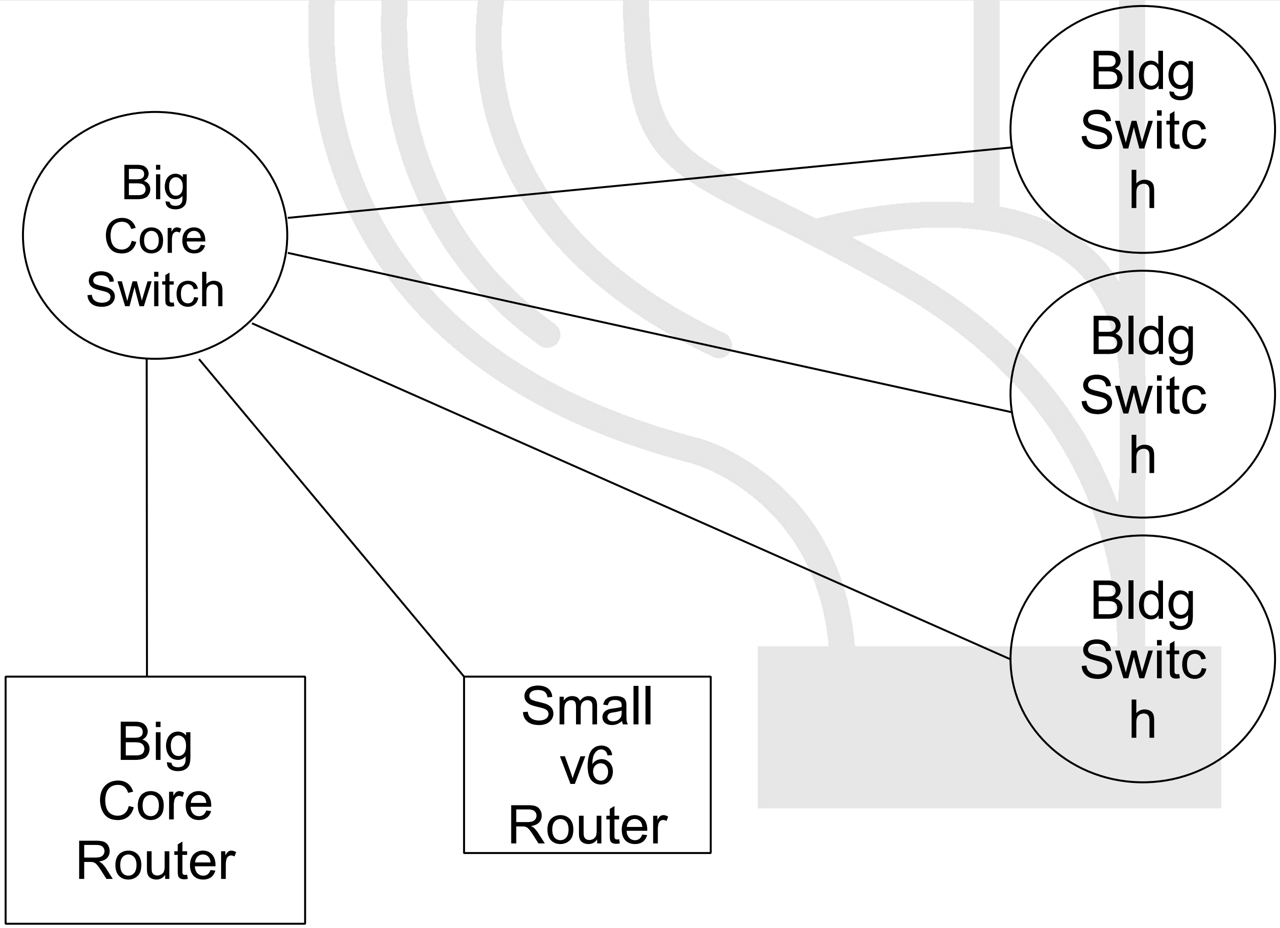

Layer-2 Campus - Redundant Switches

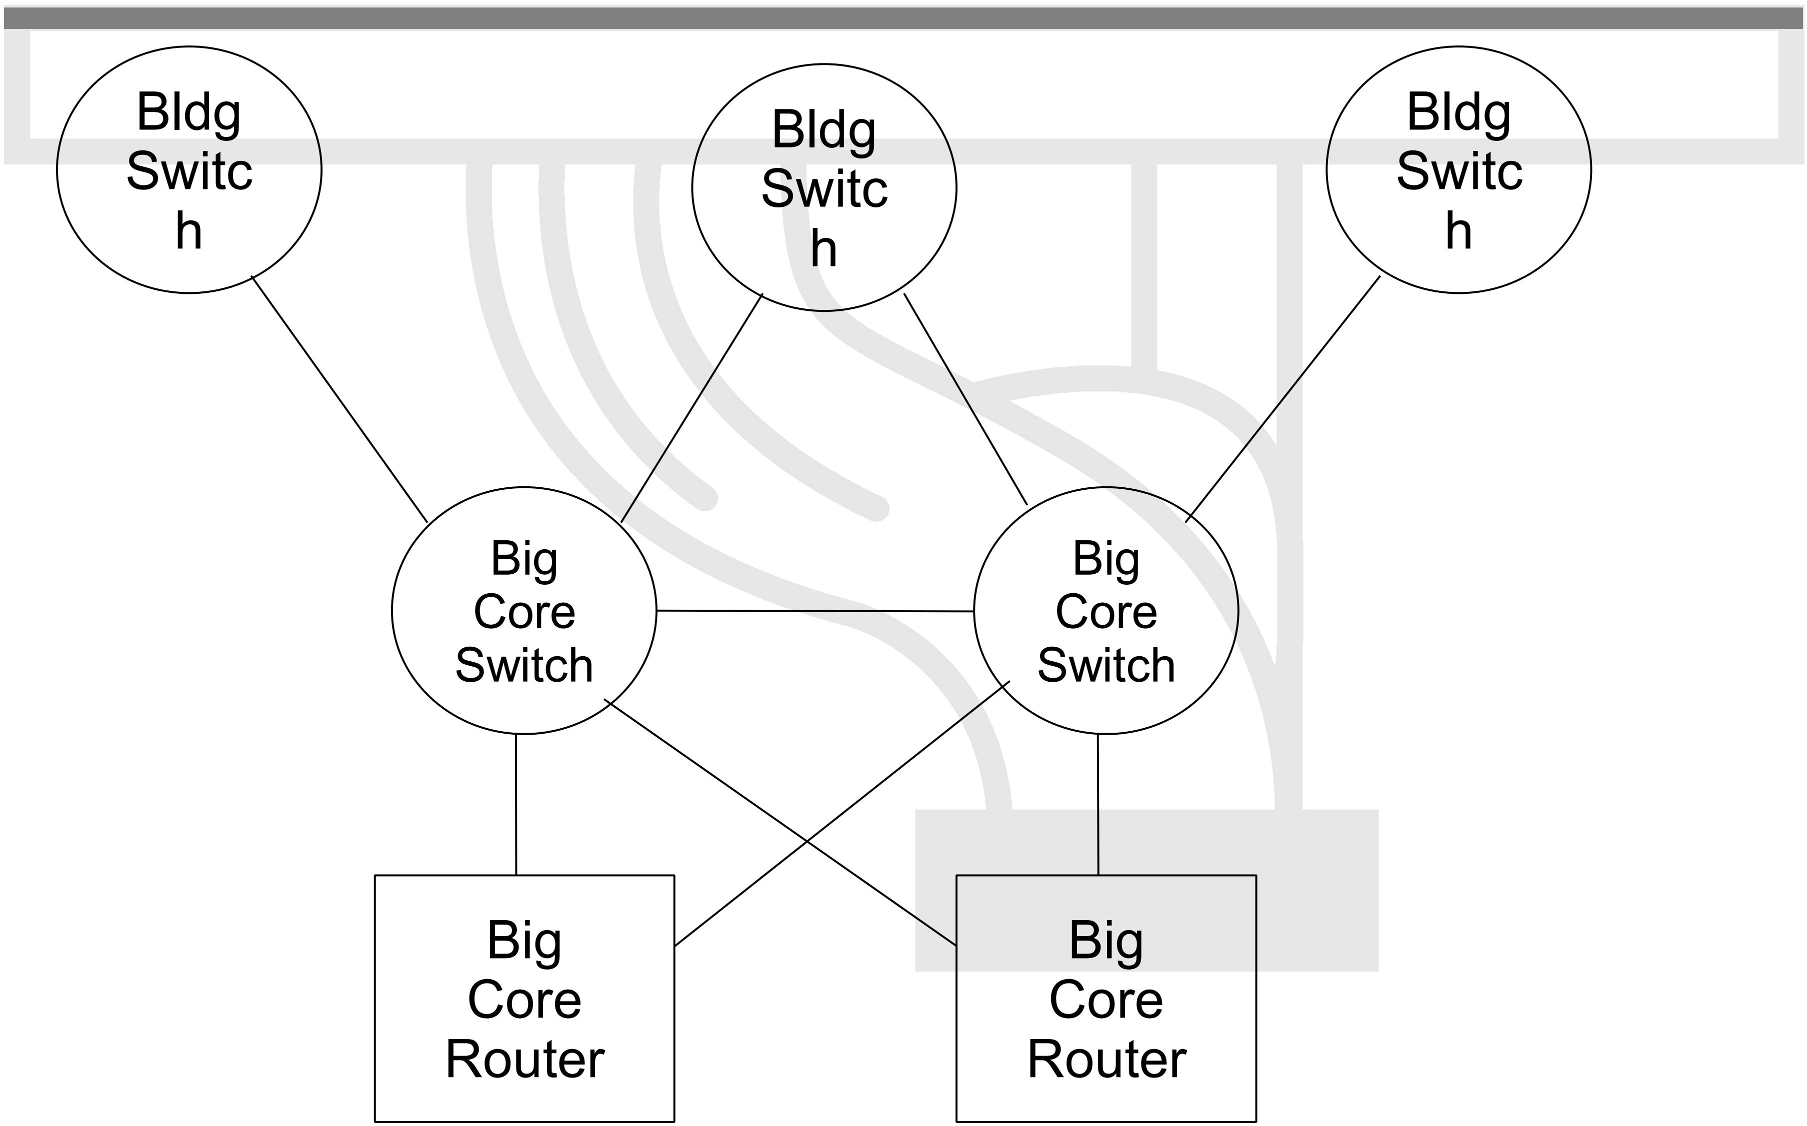

Layer-2 Campus
Redundant Switches

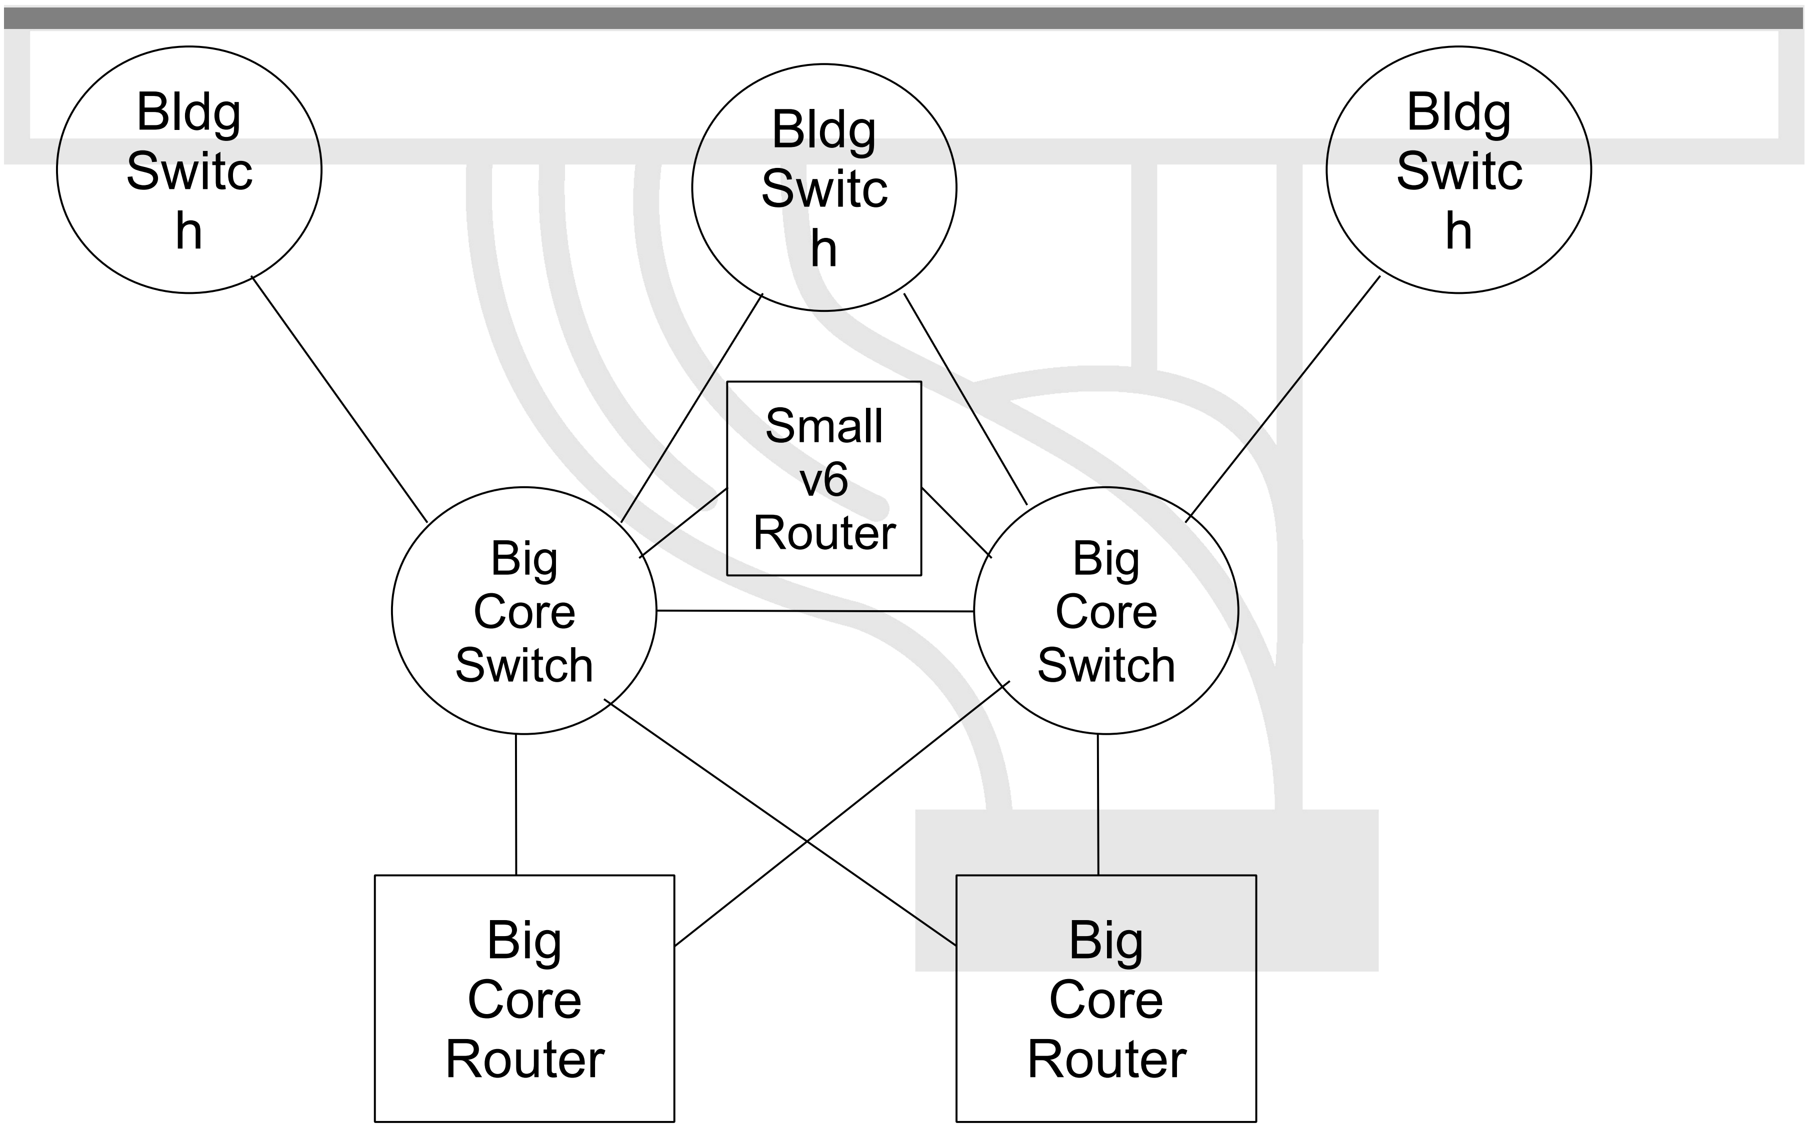

Layer-3 Campus

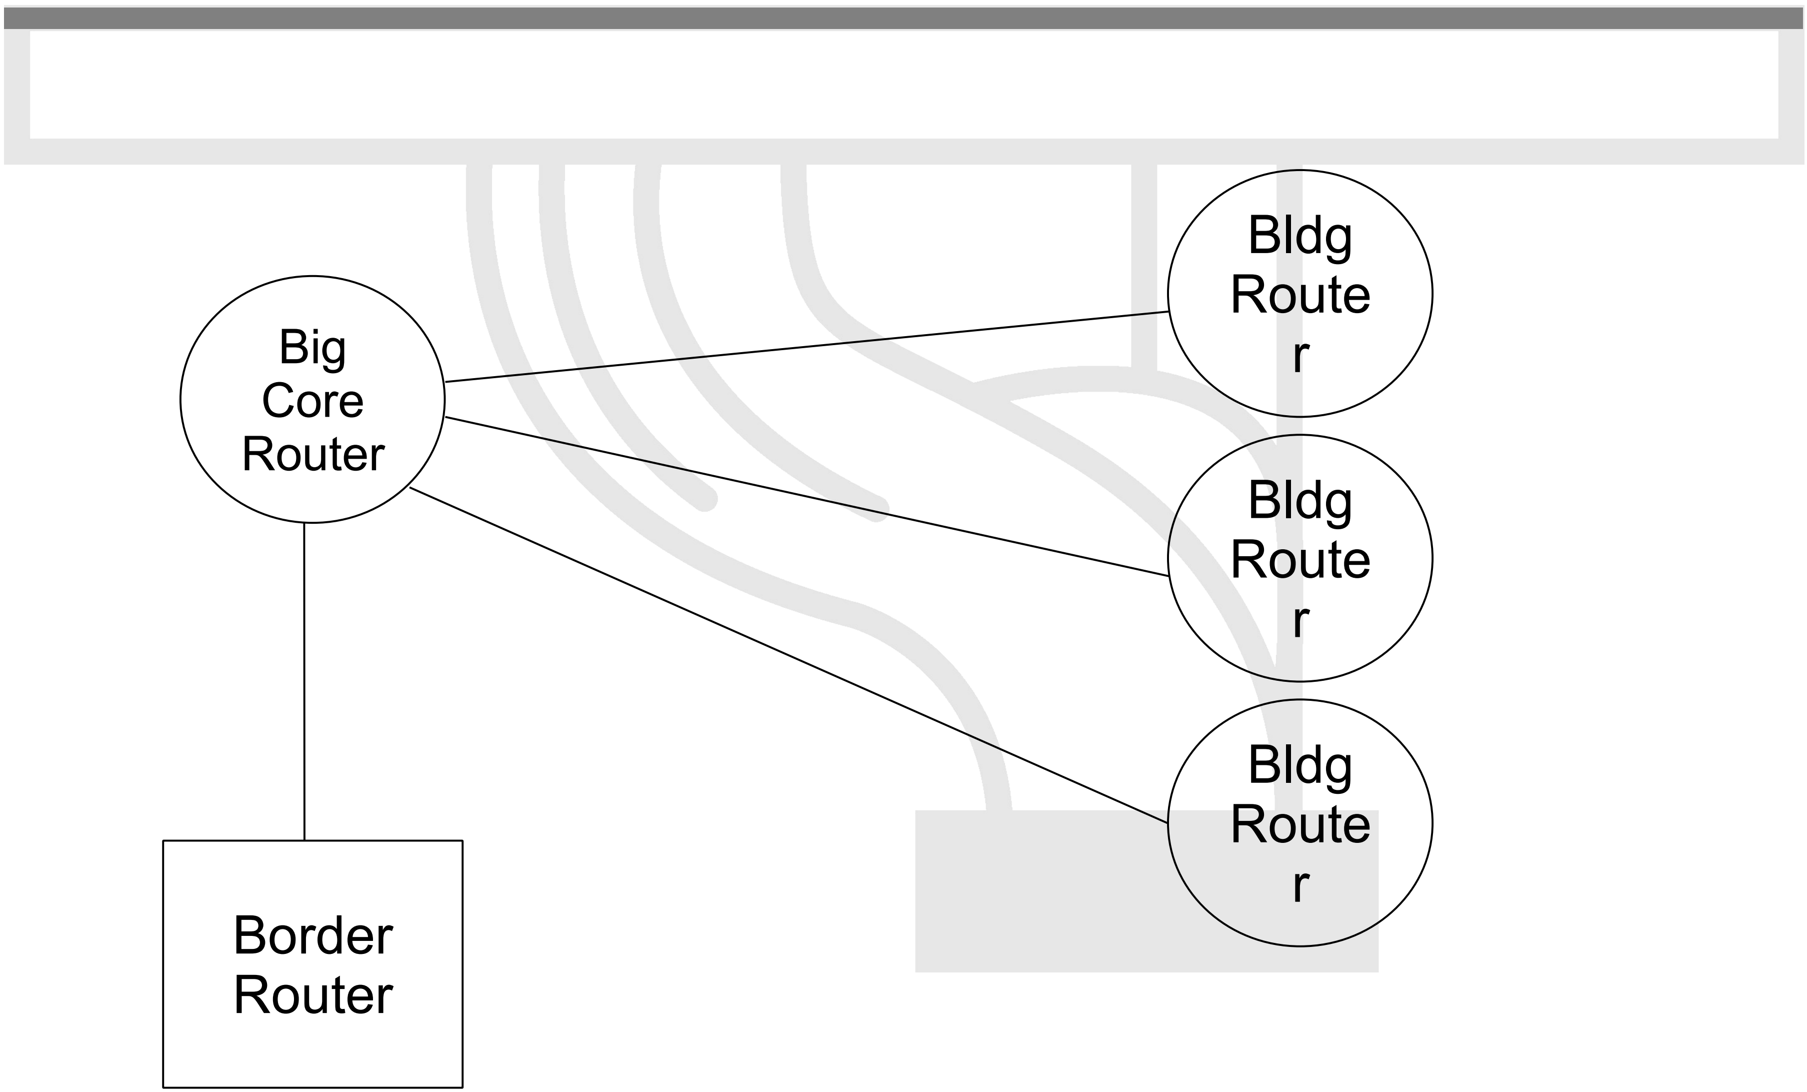

Layer-3 Campus

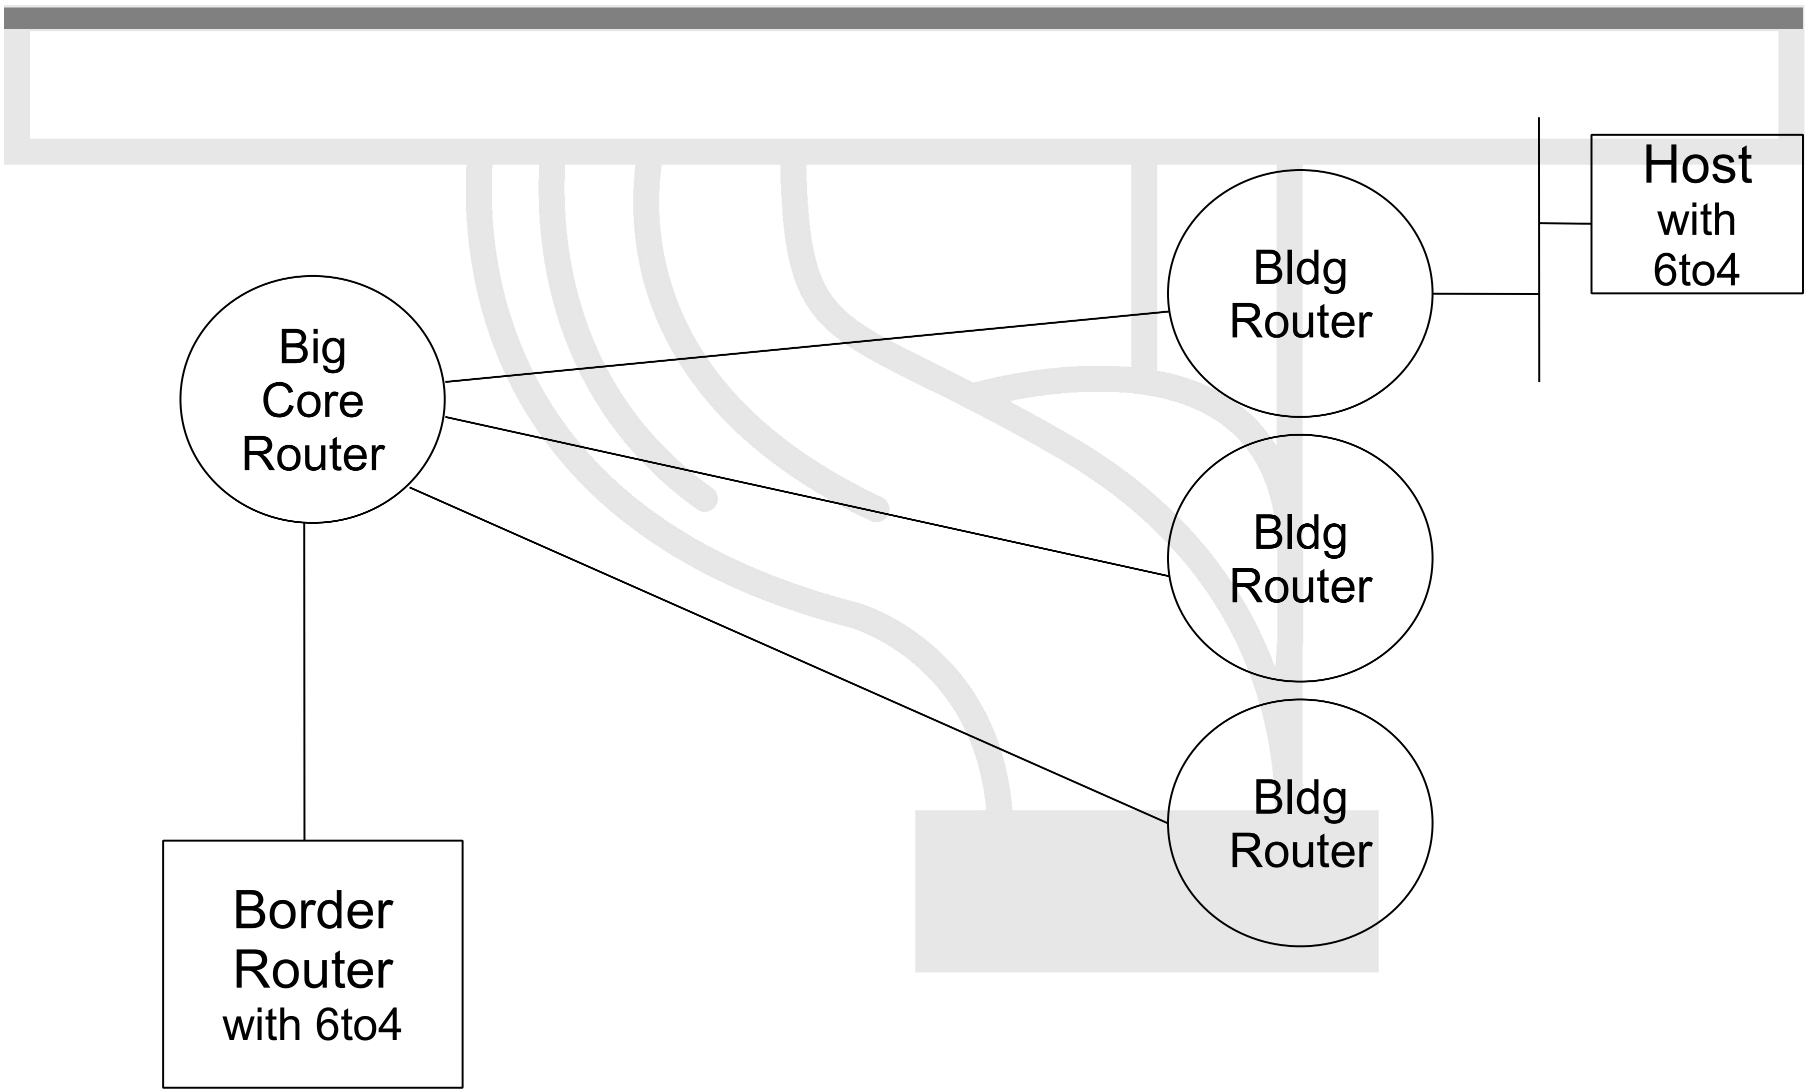

Edge Router Options

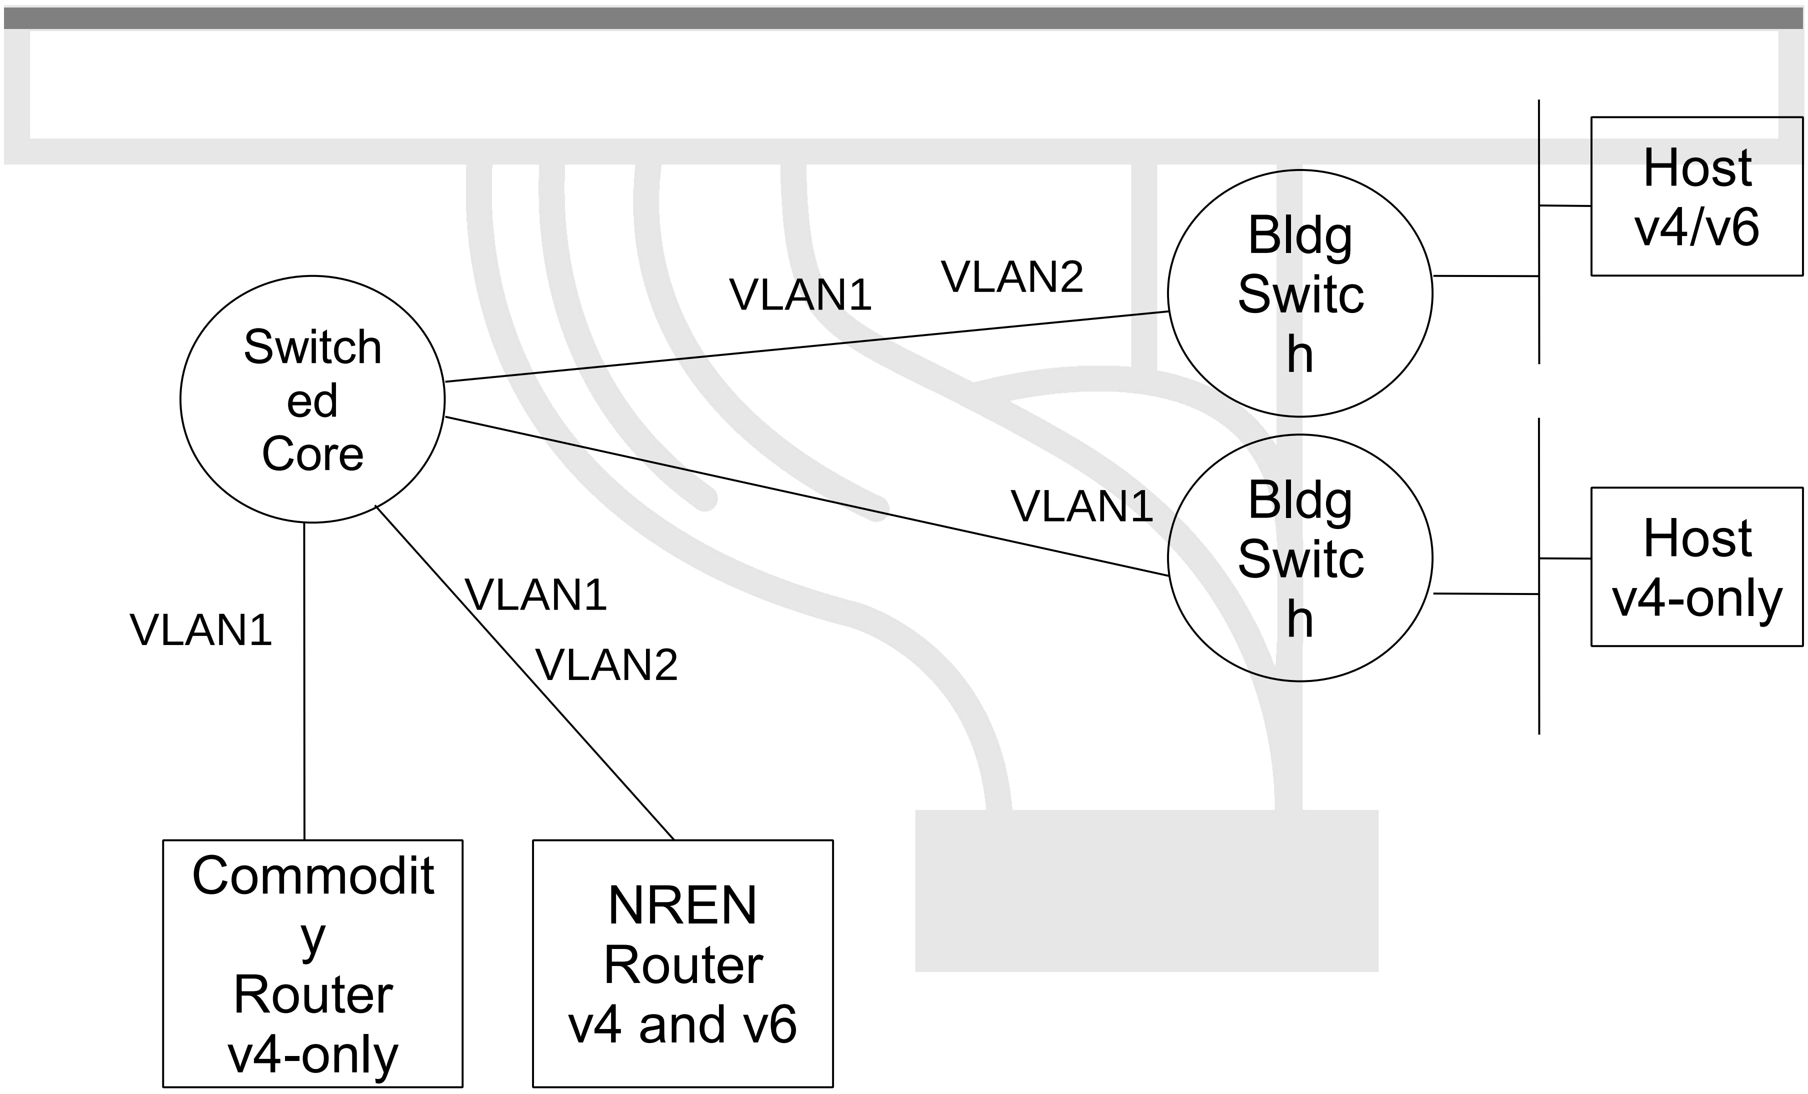

Routing Protocols

- 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.

IPv6 server configurations

Outline

- DNS
- Other applications
- Overcome IPv6 application deployment difficulties

How important is the DNS?

- Getting the IP address of the remote endpoint is necessary for every communication between TCP/IP applications
- Humans are unable to memorize millions of IP addresses (specially IPv6 addresses)
- To a larger extent: the Domain Name System (DNS) provides applications with several types of resources (domain name servers, mail exchangers, reverse lookups, …) they need
- DNS design
  - hierarchy
  - distribution
  - redundancy

DNS Lookup

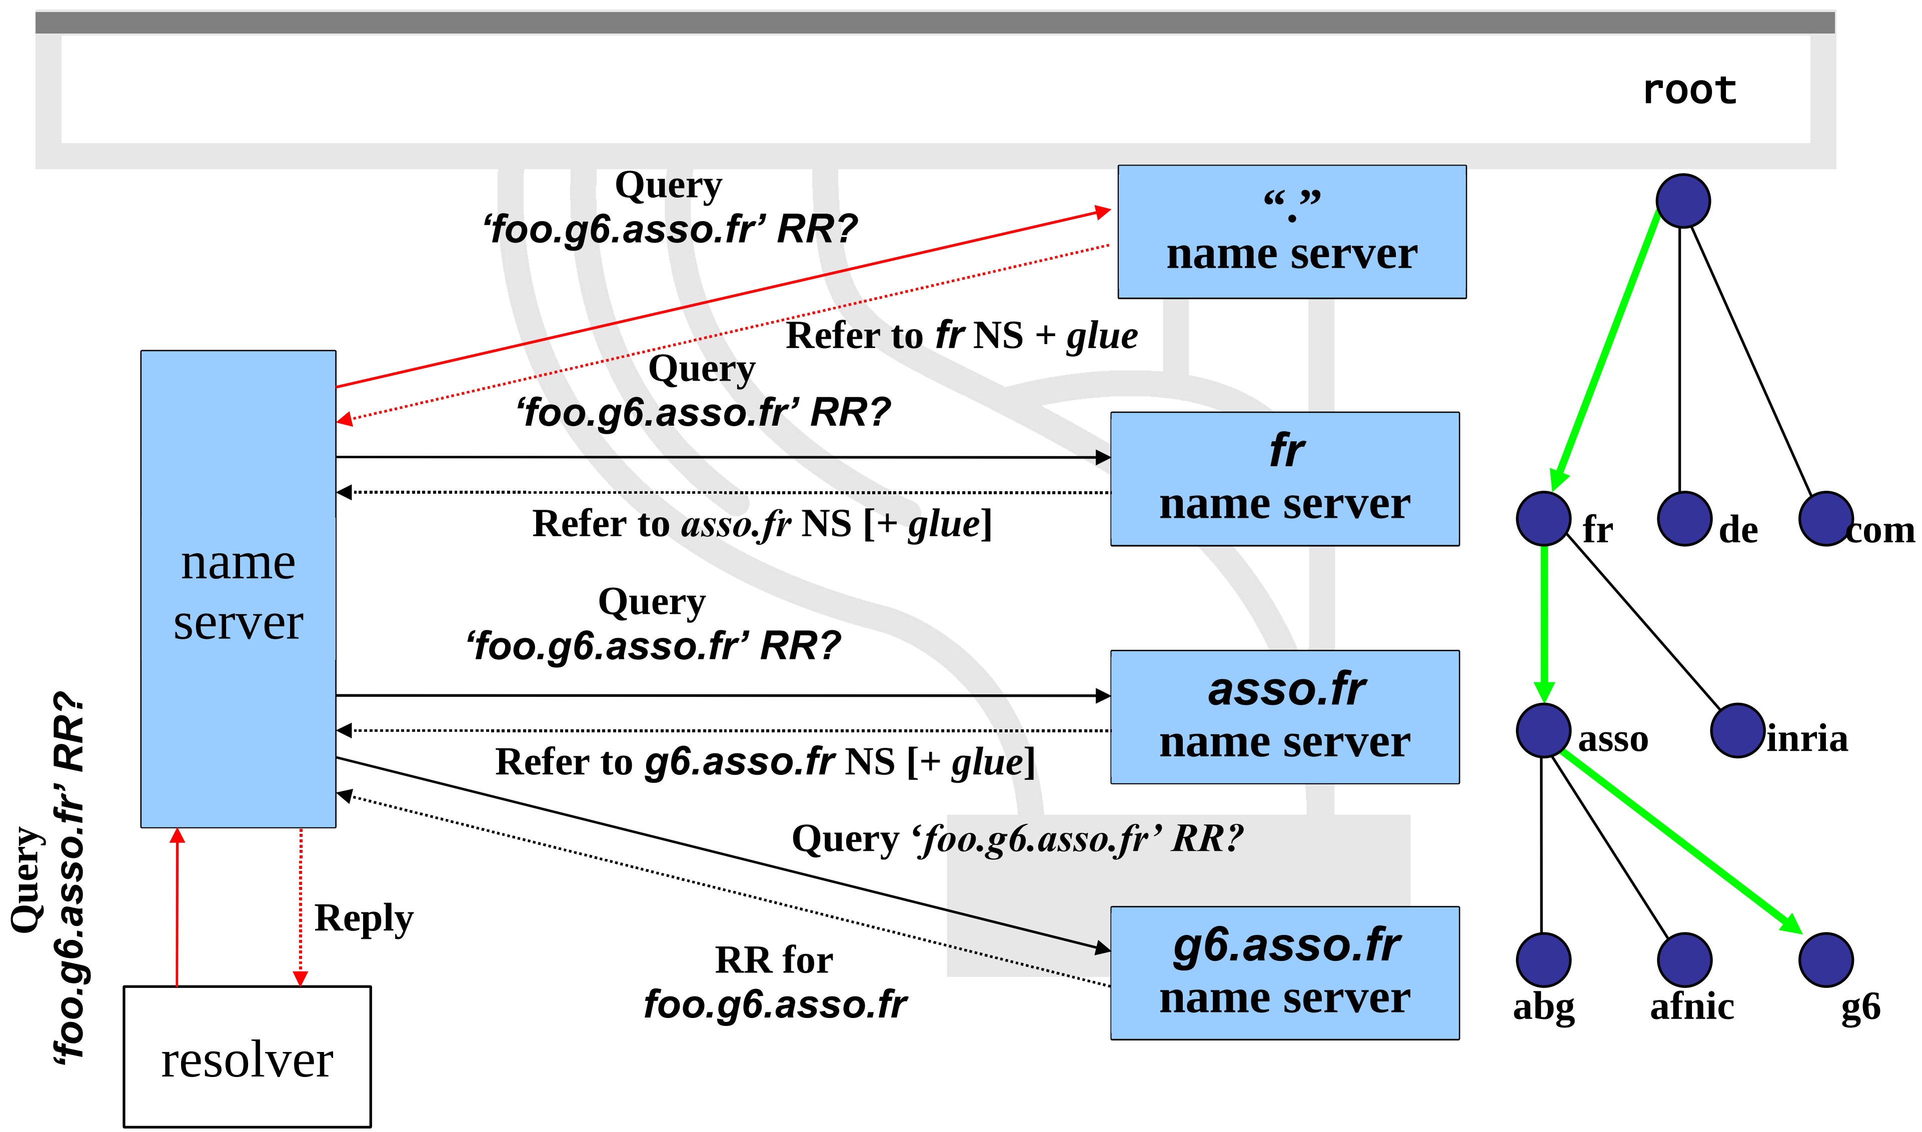

root
“.”
name server
Query
‘foo.g6.asso.fr’ RR?
Refer to fr NS + glue
Query
‘foo.g6.asso.fr’ RR?
fr
name server
Refer to asso.fr NS [+ glue]
name
server
Query
‘foo.g6.asso.fr’ RR?
asso.fr
name server
Refer to g6.asso.fr NS [+ glue]
Query ‘foo.g6.asso.fr’ RR?
Query
‘foo.g6.asso.fr’ RR?
Reply
RR for
foo.g6.asso.fr
g6.asso.fr
name server
resolver
fr
de
com
asso
inria
abg
afnic
g6

DNS Extensions for IPv6

RFC 1886 → RFC 3596 (upon successful interoperability tests)

AAAA : 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

New tree: **ip6.arpa** (under deployment)

Former tree: **ip6.int** (deprecated)

Example:

$ORIGIN 1.**0.0.0**.6.0.0.3.0.6.6.**0**.1.0.0.2.ip6.arpa.

1.**0.0.0**.1.**0.0.0.0.0.0.0.0.0.0.0** **PTR** ns3.nic.fr.

Lookups in an IPv6-aware DNS Tree

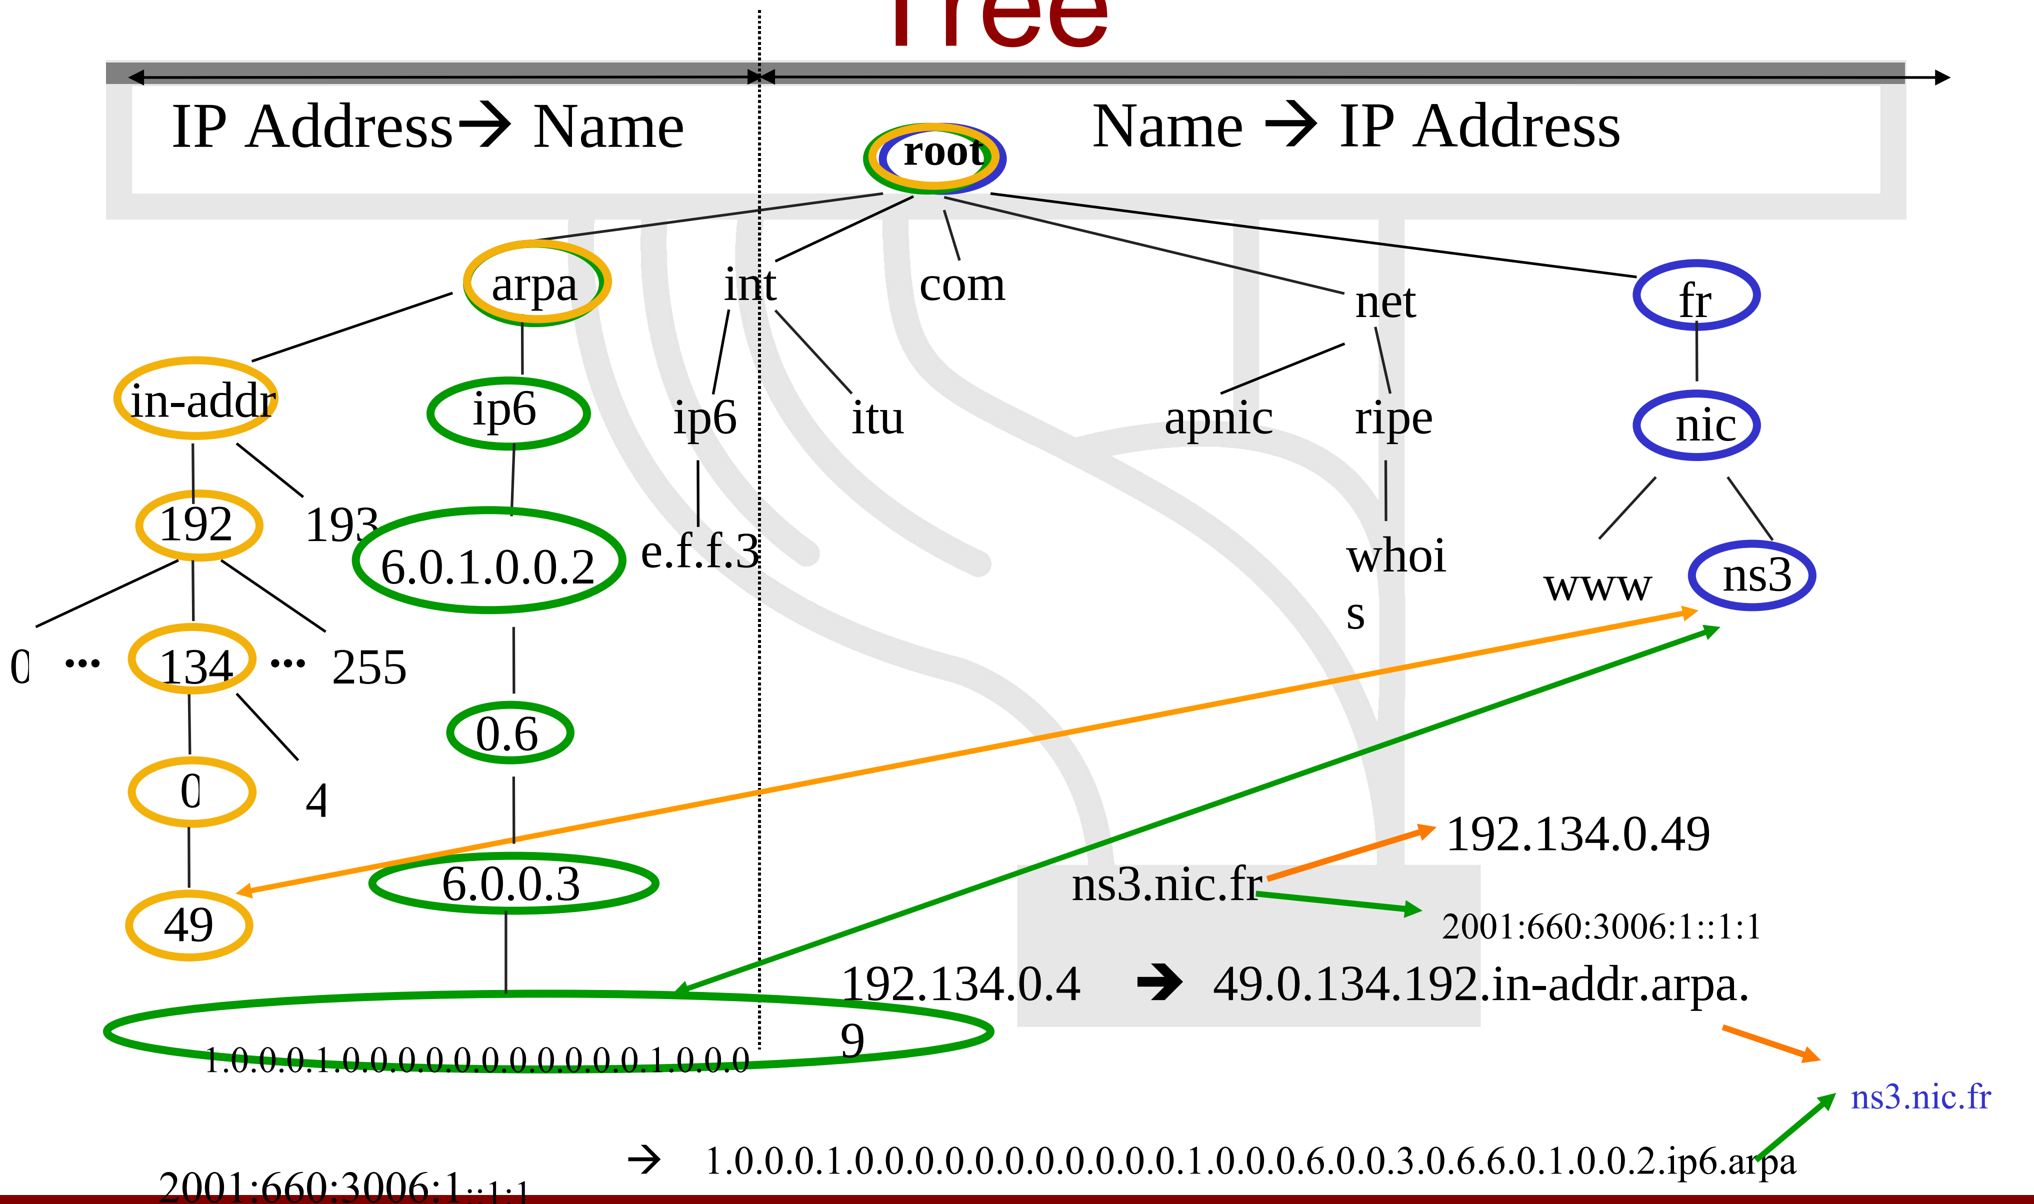

About Required IPv6 Glue in DNS Zones

When the DNS zone is delegated to a DNS server (among others) contained in the zone itself

Example: In zone file rennes.enst-bretagne.fr

```
@           IN          SOA         rsm.rennes.enst-bretagne.fr. fradin.rennes.enst-bretagne.fr.
                        (2005040201 ;serial
                        86400       ;refresh
                        3600        ;retry
                        3600000     ;expire}

                        IN          NS          rsm
                        IN          NS          univers.enst-bretagne.fr.
[…]
ipv6        IN          NS          rhadamanthe.ipv6
            IN          NS          ns3.nic.fr.
            IN          NS          rsm
;
rhadamanthe.ipv6                    IN          A           192.108.119.134
                                    IN          AAAA        2001:660:7301:1::1
[…]
```

IPv4 glue (A 192.108.119.134) is required to reach rhadamanthe over IPv4 transport
IPv6 glue (AAAA 2001:660:7301:1::1) is required to reach rhadamanthe over IPv6 transport

IPv6 DNS and root servers

- DNS root servers are critical resources!
- 13 roots « around » the world (#10 in the US)
- Not all the 13 servers already have IPv6 enabled and globally reachable via IPv6.
- Need for (mirror) root servers to be installed in other locations (EU, Asia, Africa, …)
- New technique : anycast DNS server
  - To build a clone from the master/primary server
  - Containing the same information (files)
  - Using the same IP address
- Such anycast servers have already begun to be installed :
  - F root server: Ottawa, Paris(Renater), Hongkong, Lisbon (FCCN)…
  - Look at http://www.root-servers.org for the complete and updated list.

The Two Approaches to the DNS

- The DNS seen as a Database
  - Stores different types of Resource Records (RR): SOA, NS, A, AAAA, MX, SRV, PTR, …

☐ DNS data is independent of the IP version (v4/v6) the DNS server is running on!

- The DNS seen as a TCP/IP application
  - The service is accessible in either transport modes (UDP/TCP) and over either IP versions (v4/v6)

☐ Information given over both IP versions MUST BE CONSISTENT!

DNS IPv6-capable software

- BIND (Resolver & Server)
  - http://www.isc.org/products/BIND/
  - BIND 9 (avoid older versions)
- On Unix distributions
  - Resolver Library (+ (adapted) BIND)
- NSD (authoritative server only)
  - http://www.nlnetlabs.nl/nsd/
- Microsoft Windows (Resolver & Server)

…

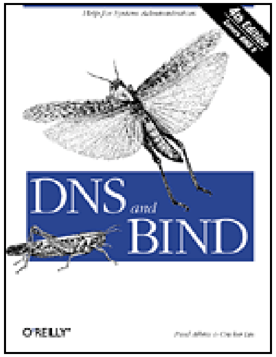

IPv6 DNS support

- BIND8
  - IPv6 RRs - only AAAA
  - IPv4 transport (IPv6 transport with patch or since 8.4.0, resolver since 8.3.0)
- BIND9
  - All IPv6 RRs
  - IPv4/IPv6 transport
- NSD
  - only authorative
- PowerDNS – SQL backend
- djbdns
  - IPv6 RRs - only AAAA
  - IPv4 transport only (IPv6 transport with patch)

Bind 9 configuration/1

- named.conf entries
  - More than one listen-on-v6 option can be used:

```
options {
        listen-on-v6 port   53 { any; };
        listen-on-v6 port 1234 { any; };
};
```

  - In order the DNS server not to server IPv6 requests. (Before 9.2.0 – now it is the default):

```
options {
        listen-on-v6 { none; };
};
```

Bind9 configuration/2

- Zone transfer:

```
transfer-source-v6 1:2:3:4:5:6:7:8;
```

- Query over IPv6 enable:

```
query-source-v6 address * 53;
```

- Don't forget to update ACLs for IPv6 addresses!

DNSv6 Operational Requirements & Recommendations

- The target today IS NOT the transition from an IPv4-only to an IPv6-only environment
- How to get there?
  - Start by testing DNSv6 on a small network and get your own conclusion that DNSv6 is harmless, but remember:
    - **The server (host) must support IPv6**
    - **And DNS server software must support IPv6**
  - Deploy DNSv6 in an incremental fashion on existing networks
  - DO NOT BREAK something that works fine (production IPv4 DNS)!

TLDs and IPv6

- One of IANA's functions is the DNS top-level delegations
- Changes in TLDs (e.g ccTLDs) has to be approved and activated by IANA
- Introduction of IPv6-capable nameservers at ccTLDs level has to be made through IANA

TLDs and IPv6 #2

How many servers supporting a domain should carry AAAA records?
  - Usually conservative approaches
  - One or two servers
- Don’t use long server names. 1024 bytes limit in DNS responses
  - Some ccTLDs had to renamed their servers (same philosophy used by root servers)

TLDs and IPv6 #3

- 17/04/2005
  - 4 TLDs (.AEROS, .NET, .COM, .INT)
  - 42 ccTLDs
- European: About half already glued
- Servers: 35 different ones, worldwide

Applications/1

- Apache
  - 2.0.x version supports IPv6 automatically
    - --enable-v4-mapped
  - Listen ::
    - Listen [::]:80
  - NameVirtualHost (IPv6 address also)
  - Access control is working – Do not forget update ACLs for IPv6 addresses
  - For Apache 1.3.14-1.3.19- there is IPv6 patch
- OpenSSH
  - ListenAddress ::
  - sshd -6 (-4)

Applications/2

- *Postfix*
  - *Postfix 2.2 officially supports IPv6*
  - *IPv6 patch and Ipv6+TLS patch for Postfix 2.1: http://www.ipnet6.org/postfix/*
  - *inet_interfaces = loopback-only" for version independent*
  - */etc/postfix/main.cf:*

```
inet_protocols = ipv4,ipv6,all
```

  - *mynetworks [ipv6:addr:range]/plen*
  - *smtp_bind_address6 Source address for outgoing SMTP connections.*
  - *lmtp_bind_address6 Source address for LMTP client connections*
- *Exim*
    - *HAVE_IPV6=YES in Local/Makefile*

Applications/3

- Sendmail
  - M4 configuration file should include IPv6 transport.
  - DAEMON_OPTIONS(‘Name=MTA-v4, Family=inet’)
  - DAEMON_OPTIONS(‘Name=MTA-v6, Family=inet6’)
  - DBMs:
    - IPv6:2002:c0a8:51d2::23f4 REJECT
  - Option:
    - ResolverOptions=WorkAroundBrokenAAAA
- No problem with having MXes with IPv6, but might be good to have a last resort MX with IPv4-only in case of broken MTAs
  - See RFC 3974

Applications/4

- Inetd
  - tcp → tcp6 or tcp46
  - udp → udp6 or udp46
- INN
  - --enable-ipv6 should be added to configure
- Diablo news server – supports IPv6
- FTP
  - vsftpd,moftpd, pure-ftpd, tnftpd, wzdftpd, lukemftpd – supports IPv6

More applications

- OpenLDAP
  - IPv6 enabled LDAP server and clients
    - Other LDAP application becomes IPv6 enabled when using OpenLDAP client libraries
  - There is also Sun ONE Directory server with IPv6
- GnomeMeeting
  - H.323 VoIP and videoconferencing. Supports IPv6 and runs on at least Linux. http://www.gnomemeeting.org/
- Kphone
  - IPv6 enabled VoIP SIP based softphone

  http://www.iptel.org/products/kphone/

Some programming languages

- Perl
  - Special modules like Socket6 and IO::Socket::INET6
- Python 2.3.4 and later works with IPv6
  - However, Windows binaries at python.org does not support it. 2.4 binaries will be built with IPv6 support
- PHP
  - Partial IPv6 support
  - Many PHP scripts work with IPv6 with no change
- Java
  - SUN Java SDK 1.4 has IPv6 support
  - Many Java applications work with IPv6 with no change due to the higher level API

IPv6 application pointers

- Very good list of applications
  http://www.deepspace6.net/docs/ipv6_status_page_apps.html
- IPv6 Application and Patch Database
  - This also has searchable interface
  http://ipv6.niif.hu/ipv6_apps/
- 6NET applications
  http://apps.6net.org/WP5Apps/Applications.html

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

Monitoring and management

Management and monitoring

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

Cricket

- Cricket is a tool for storing and viewing time-series data.
- Very flexible
- Extremely Legible Graphs
- Space and Time efficient
- Platform Independent

Example Graphs

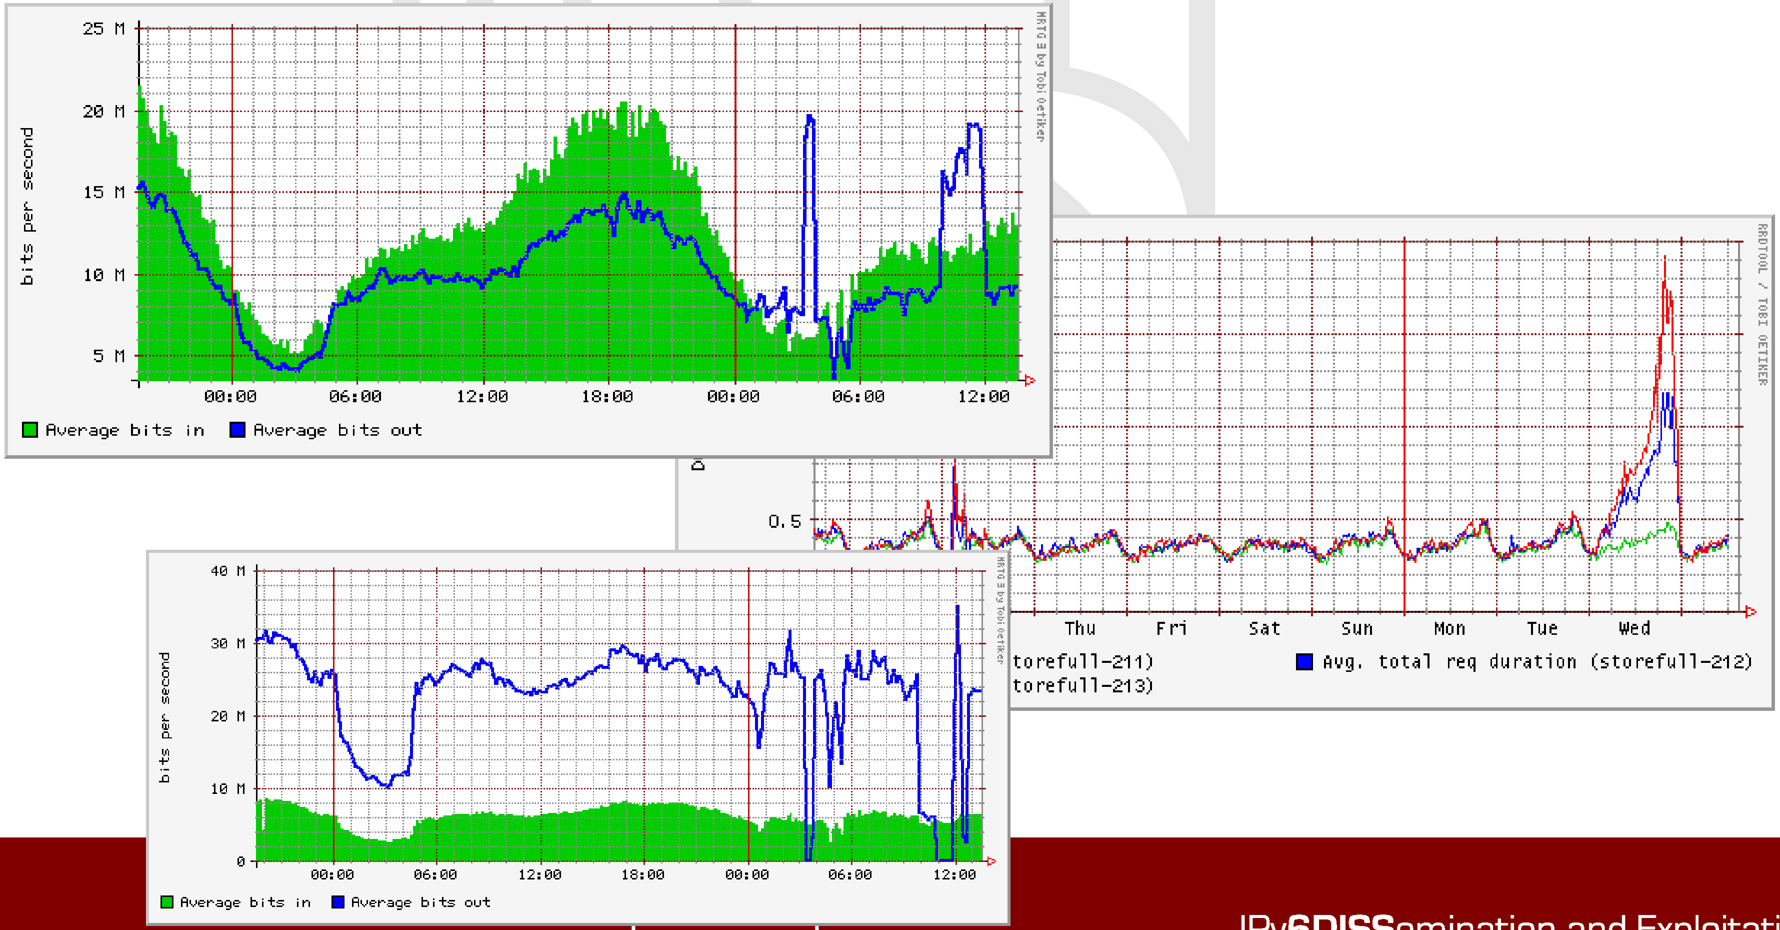

Cricket and IPv6

- No separate SNMP MIBs for IPv6 traffic implemented yet
  - Separate IPv6 infrastructure – easy to monitor
  - Dual-stack infrastructure – no easy way to monitor
    - firewall filter and counters – hardly possible on Cisco
    - From CLI: show interface accounting – misleading implementations – only process switched packets on GSR+E3 cards it is correct

Nagios: Overview

- Web-based monitoring system
- Allows for monitoring of virtually any service the NOC provides
- Provides alerting and acknowledgment capabilities
- Provides logging of downtimes and reporting capabilities

Interface

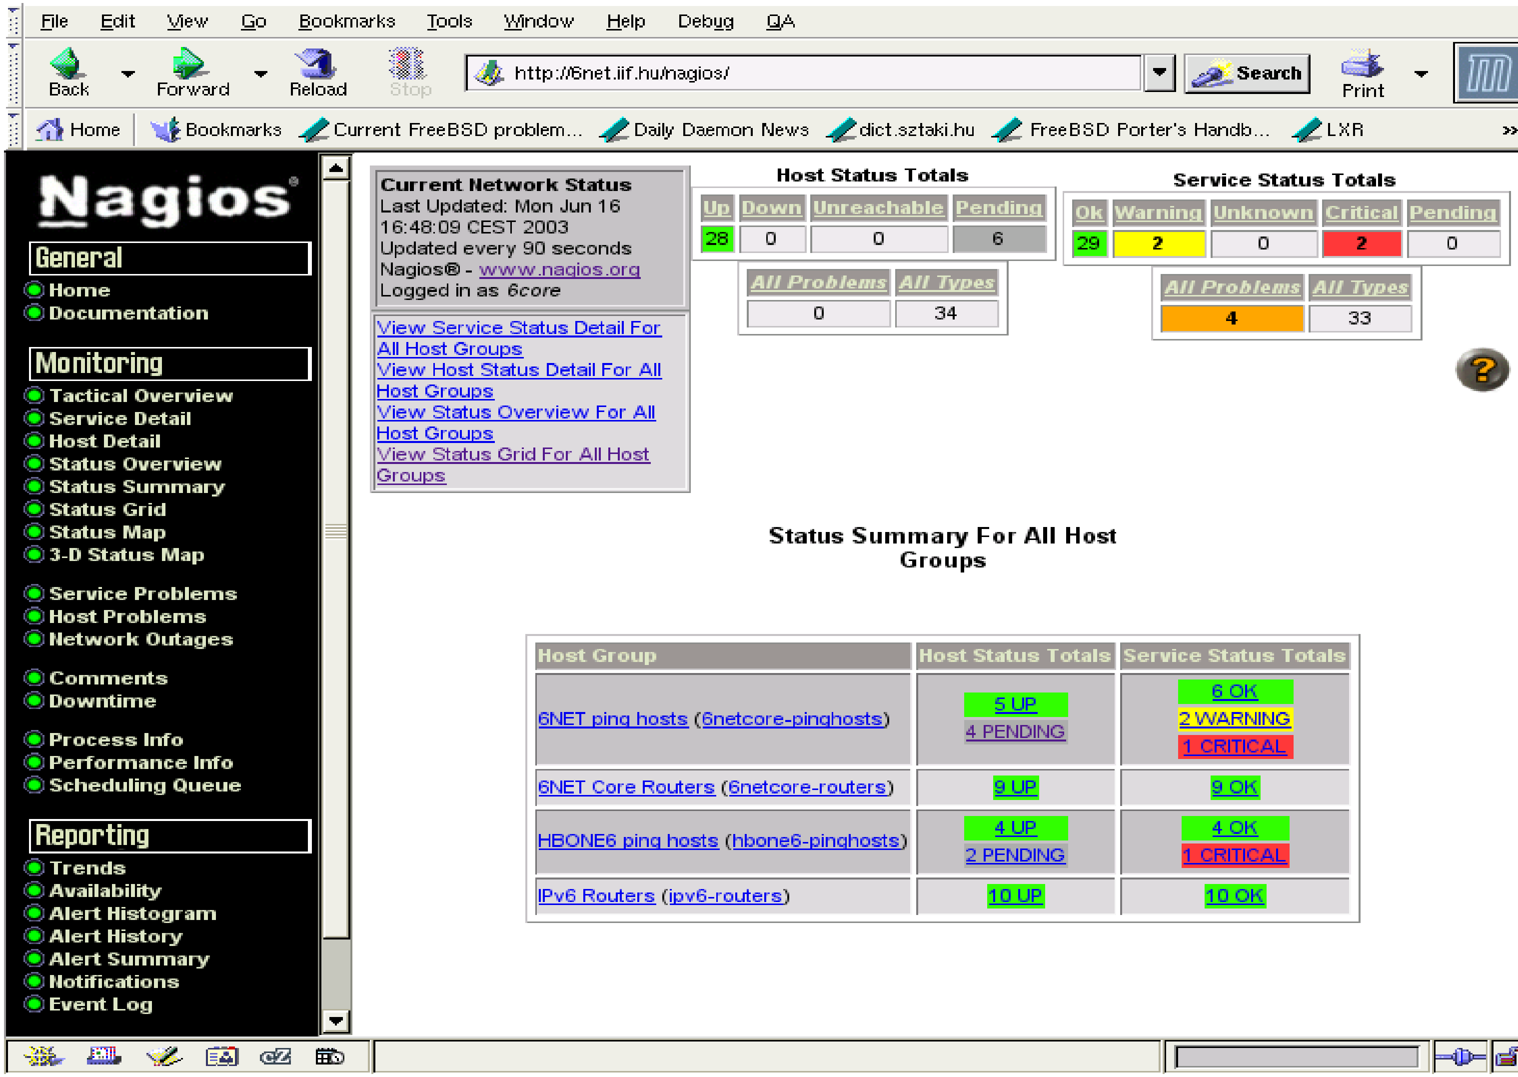

File Edit View Go Bookmarks Tools Window Help Debug QA
Back Forward Reload Stop
http://6net.iif.hu/nagios/
Search
Print
Home Bookmarks Current FreeBSD problem... Daily Daemon News dict.sztaki.hu FreeBSD Porter's Handb... LXR
Nagios
General
Home
Documentation
Monitoring
Tactical Overview
Service Detail
Host Detail
Status Overview
Status Summary
Status Grid
Status Map
3-D Status Map
Service Problems
Host Problems
Network Outages
Comments
Downtime
Process Info
Performance Info
Scheduling Queue
Reporting
Trends
Availability
Alert Histogram
Alert History
Alert Summary
Notifications
Event Log
Current Network Status
Last Updated: Mon Jun 16 16:48:09 CEST 2003
Updated every 90 seconds
Nagios® - www.nagios.org
Logged in as 6core
View Service Status Detail For All Host Groups
View Host Status Detail For All Host Groups
View Status Overview For All Host Groups
View Status Grid For All Host Groups
Host Status Totals
Up Down Unreachable Pending
28 0 0 6
All Problems All Types
0 34
Service Status Totals
Ok Warning Unknown Critical Pending
29 2 0 2 0
All Problems All Types
4 33
Status Summary For All Host Groups
Host Group
Host Status Totals
Service Status Totals
6NET ping hosts (6netcore-pinghosts)
5 UP
4 PENDING
6 OK
2 WARNING
1 CRITICAL
6NET Core Routers (6netcore-routers)
9 UP
9 OK
HBONE6 ping hosts (hbone6-pinghosts)
4 UP
2 PENDING
4 OK
1 CRITICAL
IPv6 Routers (ipv6-routers)
10 UP
10 OK

IPv6 status

- Monitoring
  - Ping over IPv6 OK – with plugin
  - TCP services over IPv6 OK – with plugin
  - UDP services over IPv6 OK – with plugin
  - SNMP over IPv6 Not yet - working on it

RANCID: Really Awesome New Cisco ConfIg Differ

- Web-based CVS repository of configuration changes
- Unix cron jobs at regular intervals check configured routers for configuration changes
- If a change is detected, RANCID e-mails all the engineers with the changes and updates the CVS repository
- Web-based CVS repository allows engineers to choose arbitrary dates to view configuration changes
- Can alter scripts to grab any information from the router that you want to track

Output of Rancid

Return to cntrl.6net.hbone.hu CVS log — Up to [6NET router configs] / 6net / configs

File: [6NET router configs] / 6net / configs / cntrl.6net.hbone.hu
Revision **1.156**: **download** - view: **text**, annotated - select for diffs - revision graph
Thu Aug 5 16:15:10 2004 UTC (5 weeks, 3 days ago) by *mohacsi*
Branches: MAIN
CVS tags: HEAD

updates

```
!RANCID-CONTENT-TYPE: cisco
!
!Chassis type: 7206VXR - a 7200 router
!CPU: NPE400, R7000 CPU at 350MHz, impl 39, Rev 3.3, 256KB L2 Cache
!
!Memory: main 491520K/32768K
!Memory: nvram 125K
!Memory: bootflash 8192K
!Memory: pcmcia ATA slot0 125952K
!
!Processor ID: 28712851
!
!Power: Power Supply 1 is Zytek AC Power Supply. Unit is on.
!Power: Power Supply 2 is Zytek AC Power Supply. Unit is on.
!
!Image: Software: C7200-P-M, 12.3(7)T1, RELEASE SOFTWARE (fc2)
```

Netflow

NetFlow for IPv6

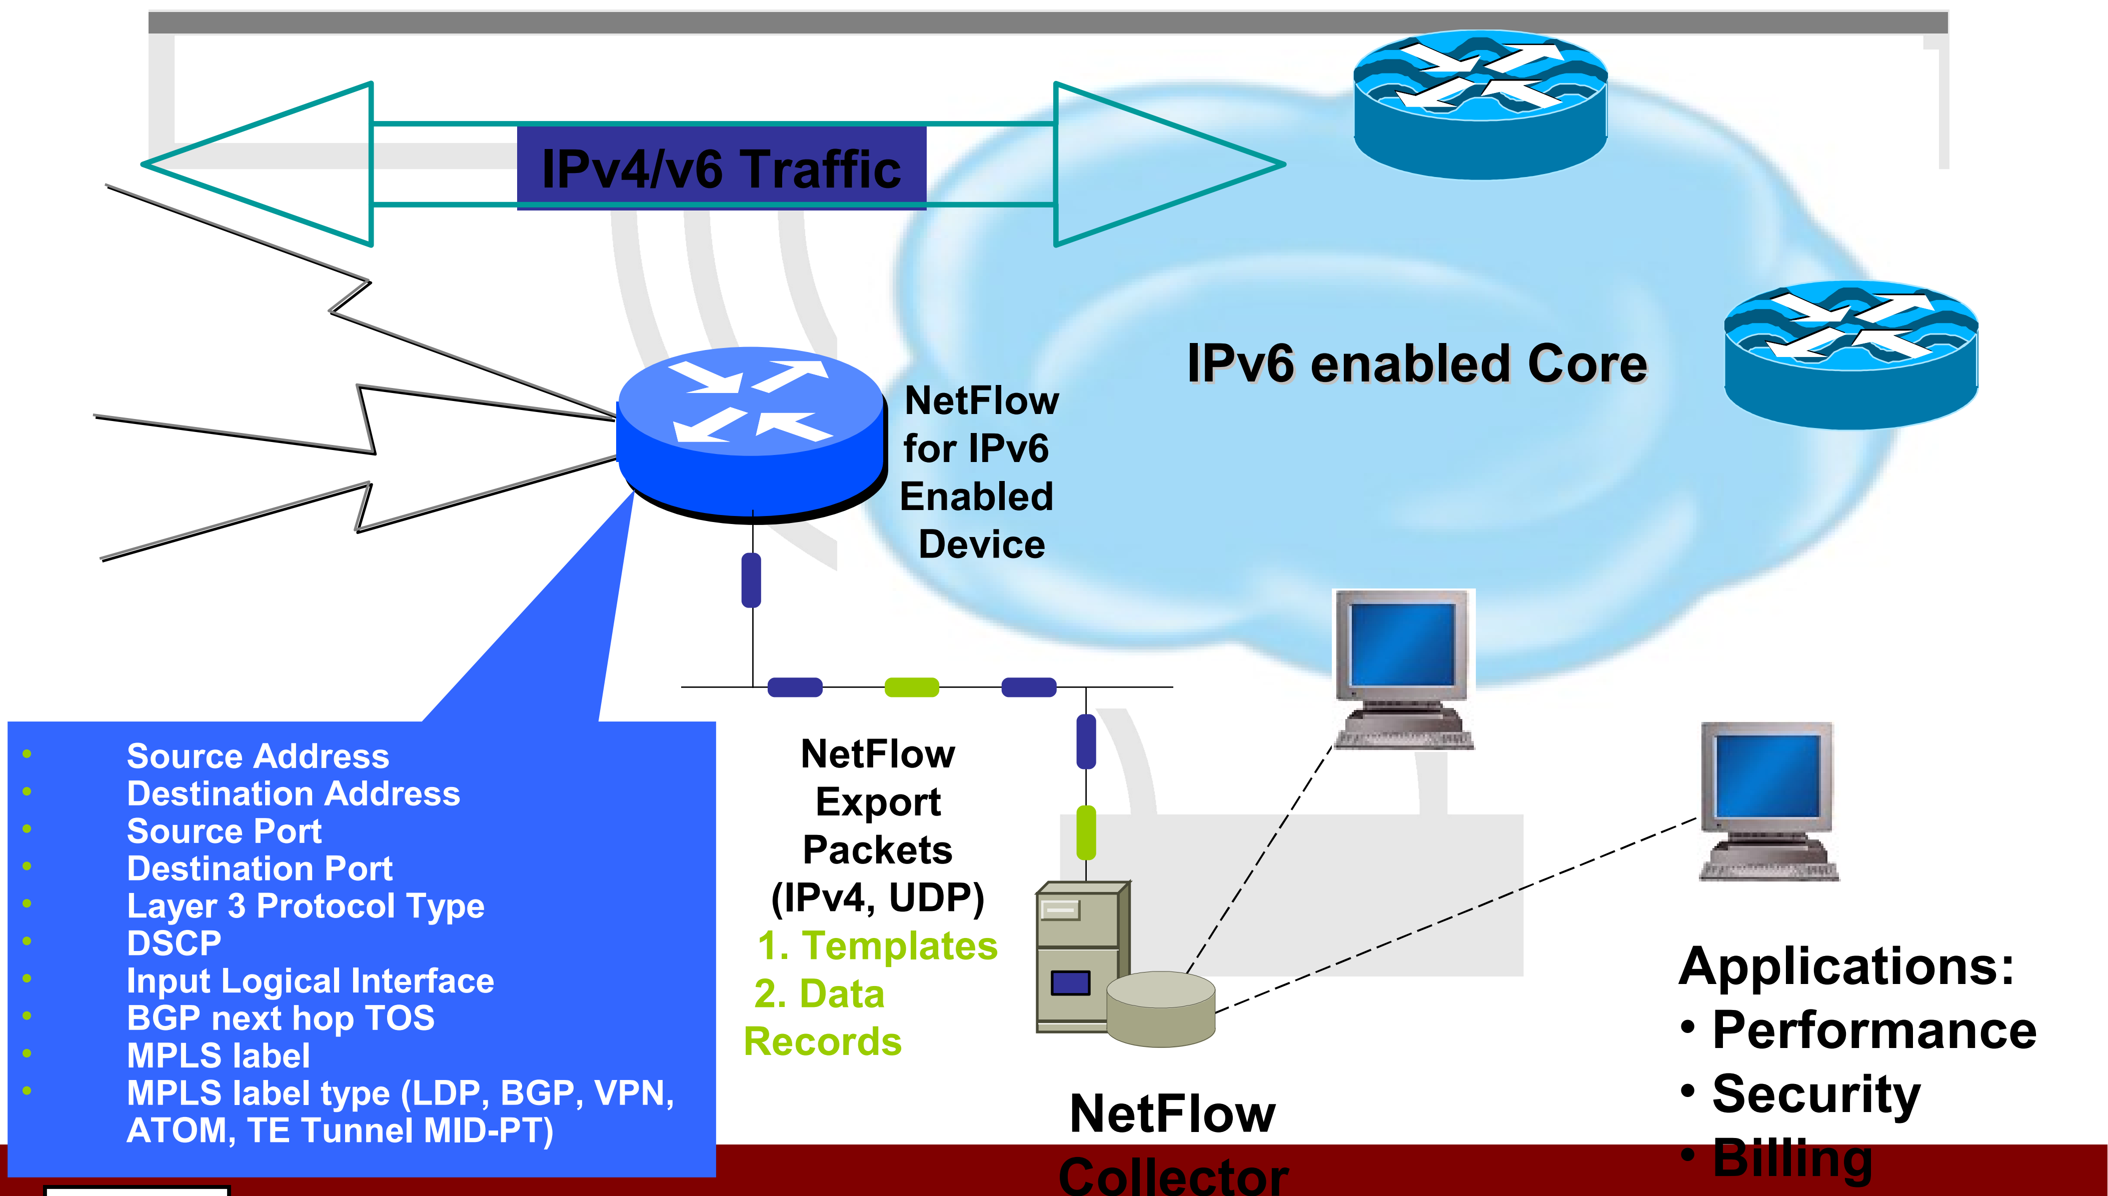

NetFlow Version 9

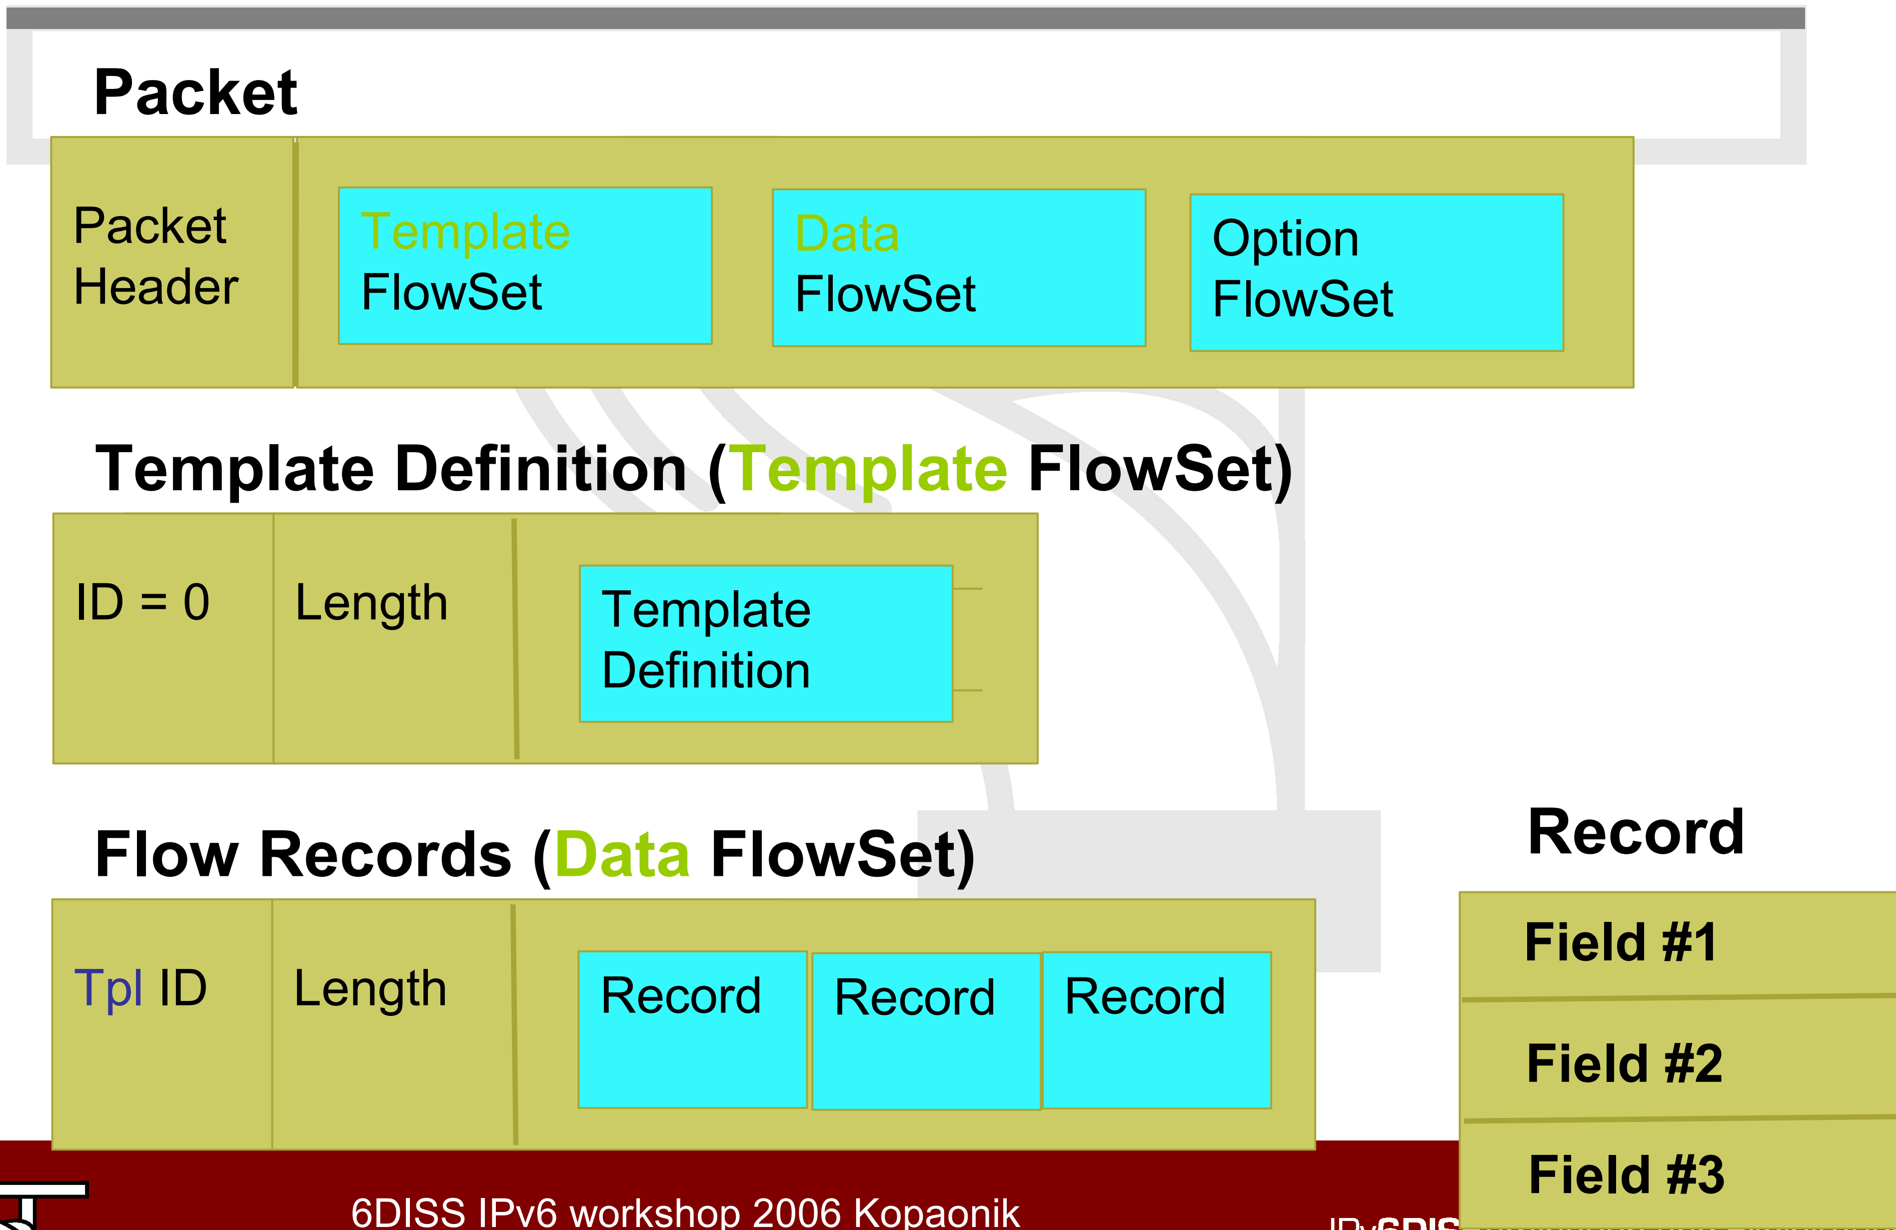

Looking Glass

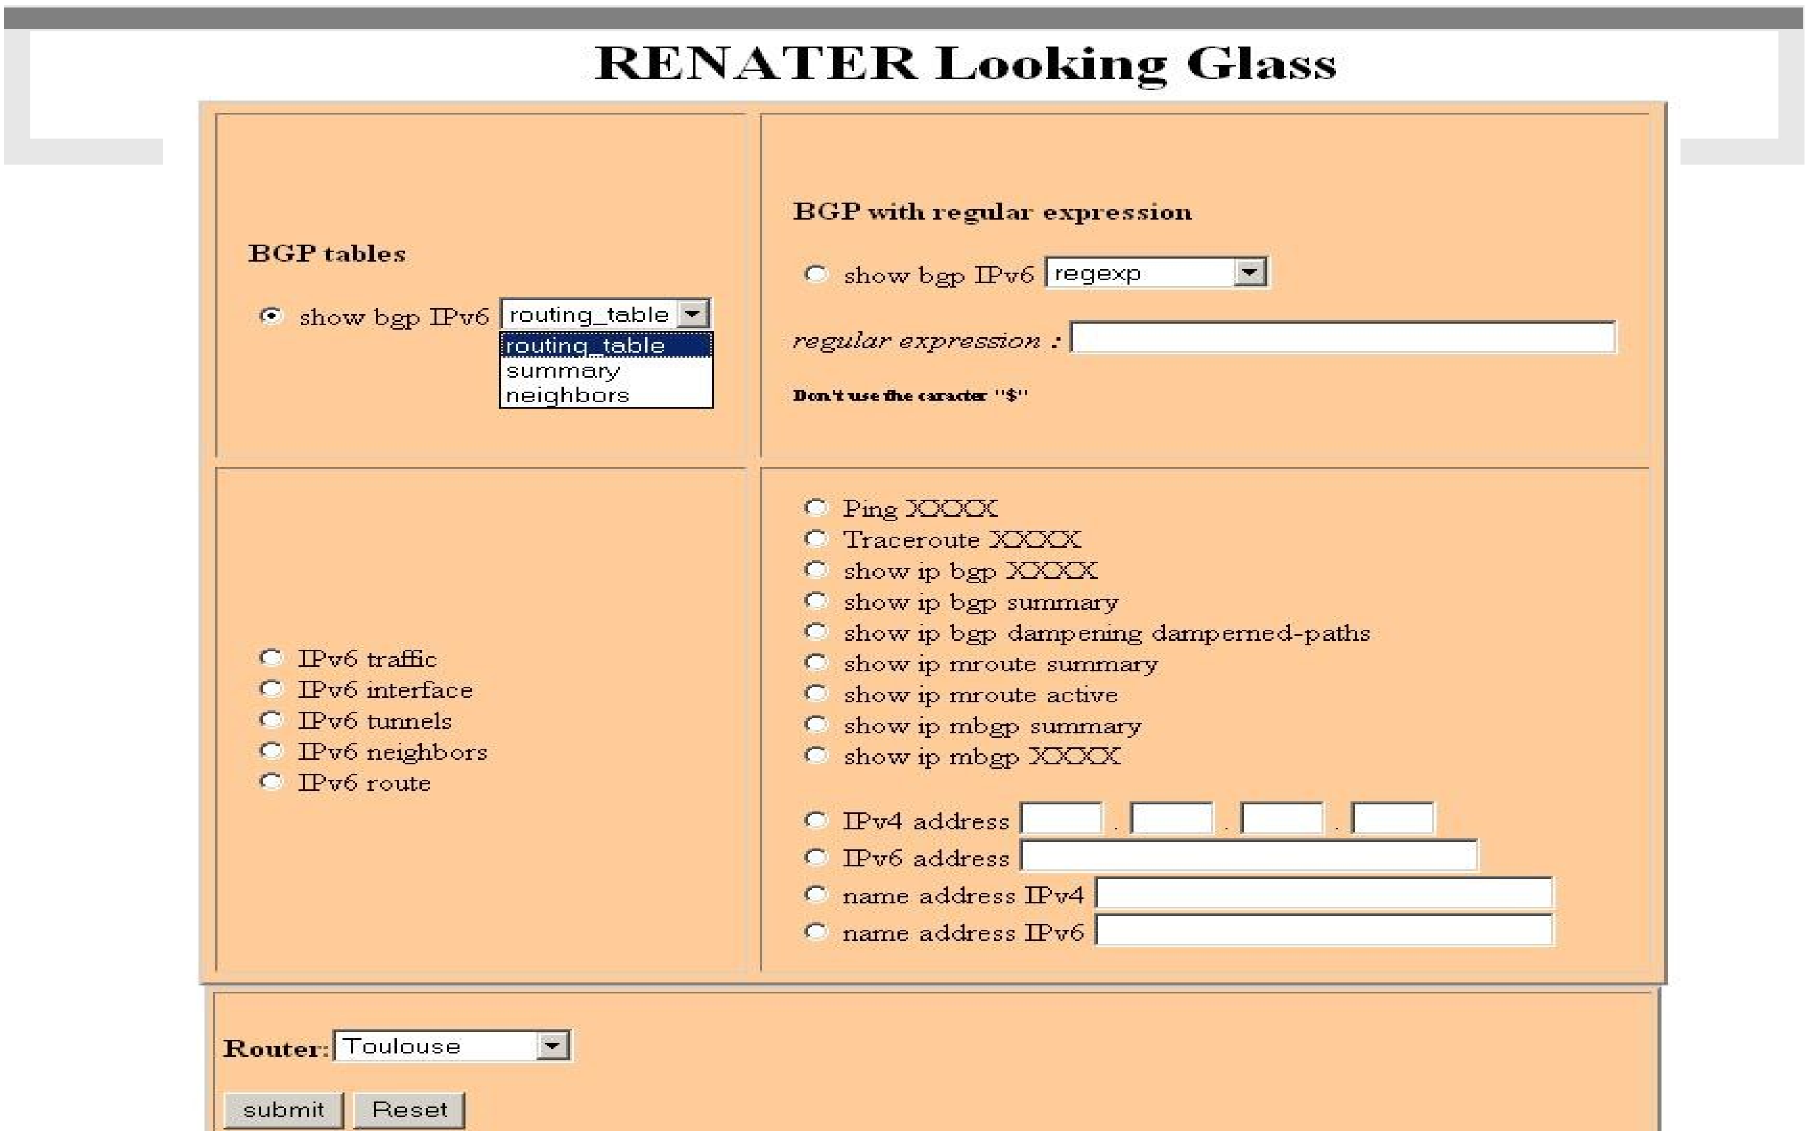

LAN IPv6 management

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