Dual Stacking a NOC

Dave Wilson, RIPE-46
DW238-RIPE

Dual stack: network and staff

- What lay before us
- Phases of deployment
- First steps
- Managed services
- Rollout

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

The problem statement

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

The IPv4 network

- Core of Cisco 12000s, one or two each in four – no, five PoPs – wait, six...
- Access networks mostly 7200/7500s, some layer 3 switches
- One Juniper M20
- Linked by GigEther/STM-1

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

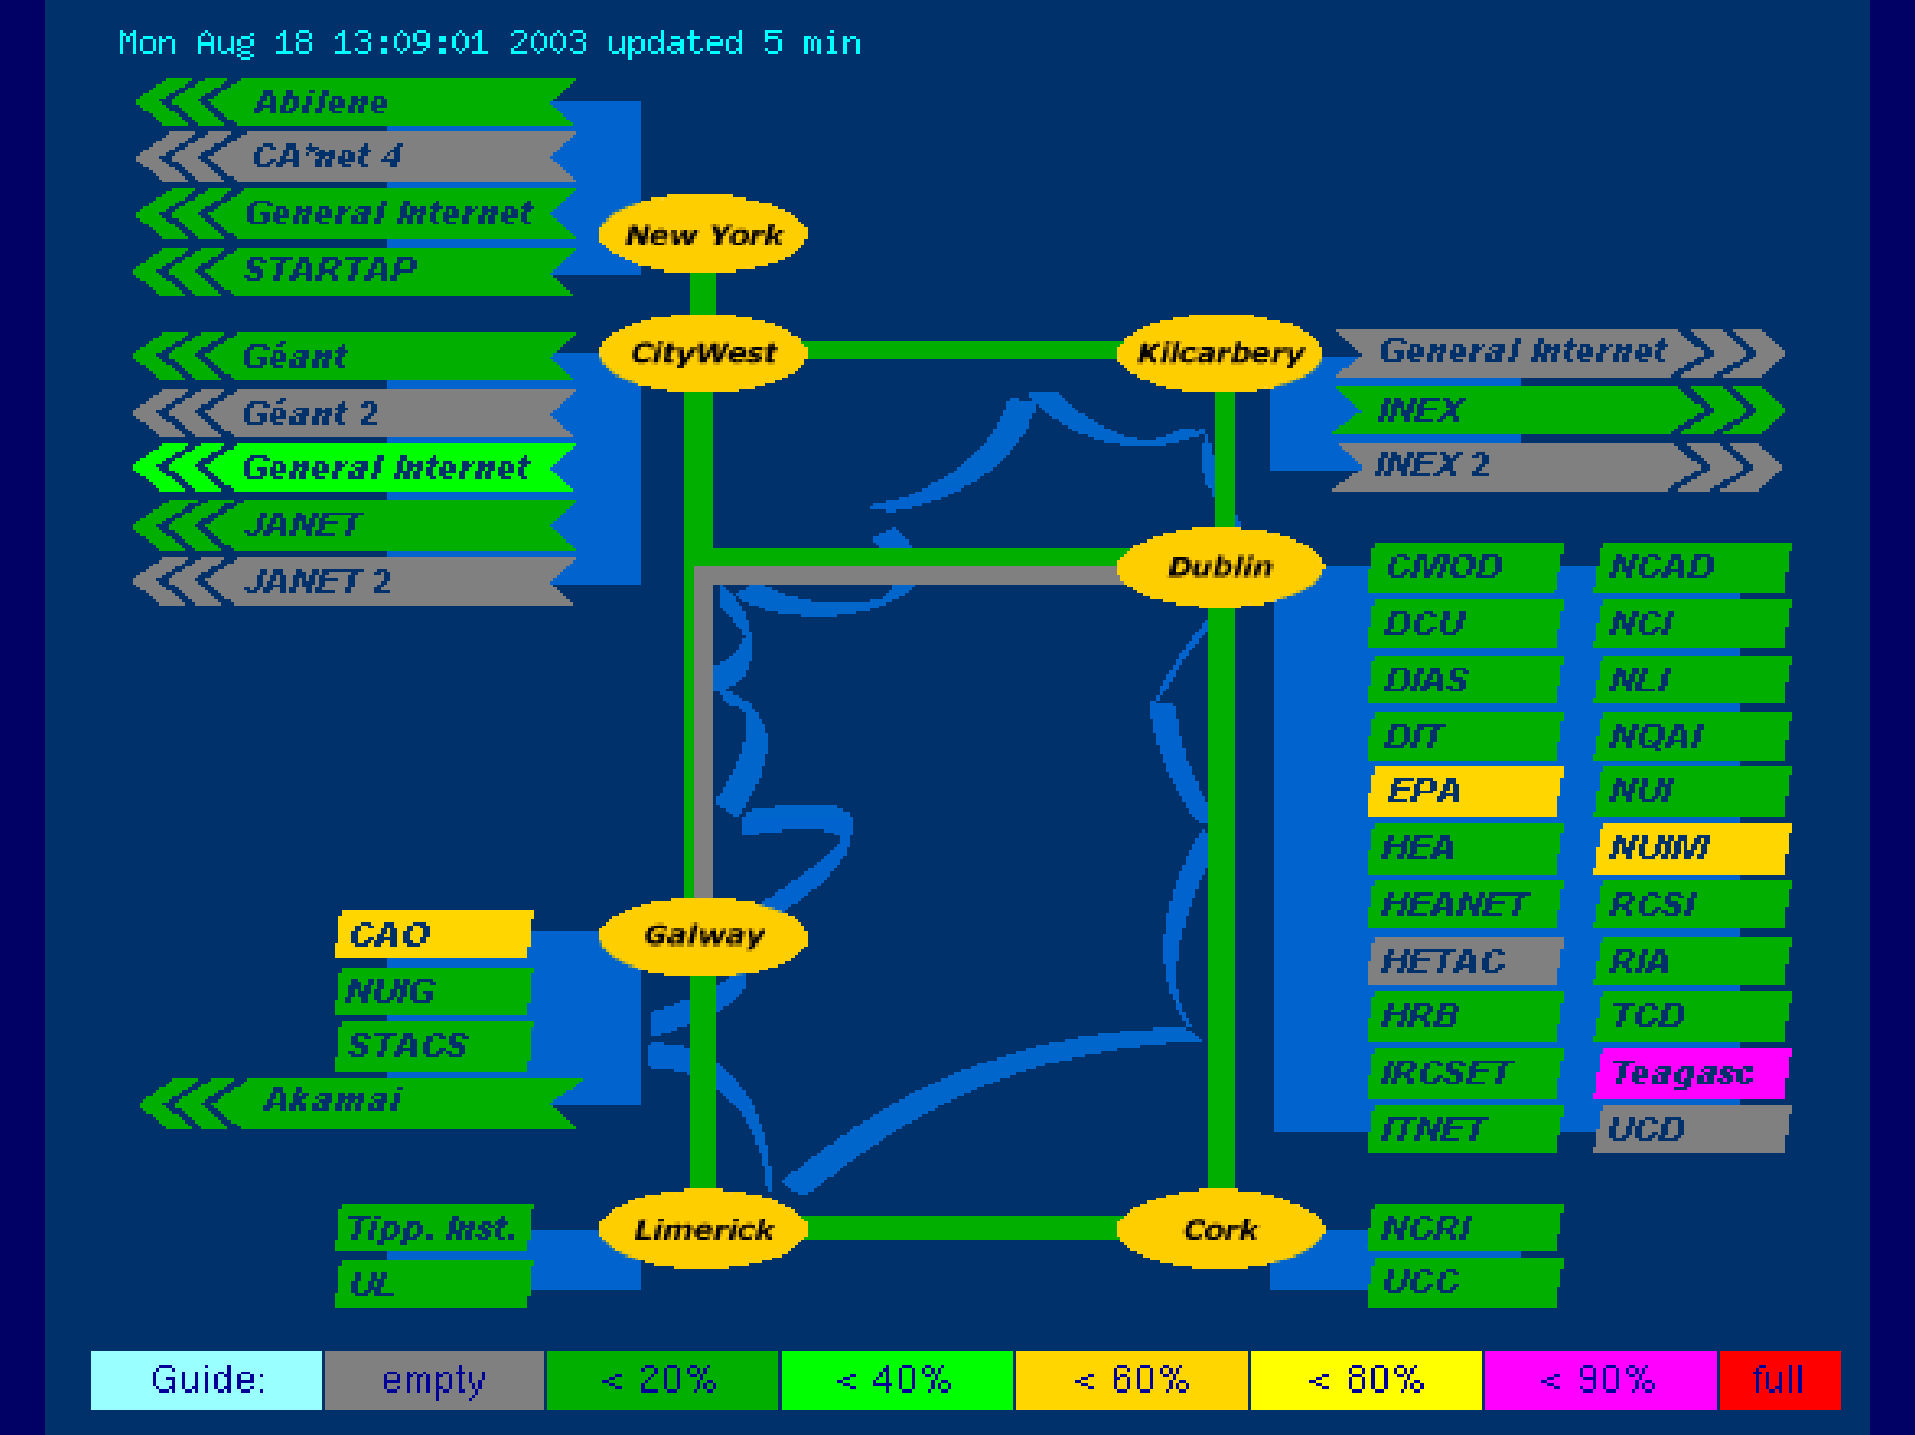
Mon Aug 18 13:09:01 2003 updated 5 min
Abilene
CA*net 4
General Internet
STARTAP
New York
Géant
Géant 2
General Internet
JANET
JANET 2
CityWest
Kilcarbery
General Internet
INEX
INEX 2
Dublin
CMOD
DCU
DIAS
DIT
EPA
HEA
HEANET
HETAC
HRB
IRCSET
ITNET
NCAD
NCI
NLI
NQAI
NUI
NUIM
RCSI
RIA
TCD
Teagasc
UCD
CAO
NUIG
STACS
Akamai
Galway
Tipp. Inst.
UL
Limerick
Cork
NCRI
UCC
Guide:
empty
< 20%
< 40%
< 60%
< 80%
< 90%
full

The IPv4 systems

- Heterogenous hardware and software
- Mostly UNIX based (Solaris, Tru64, Red Hat) and some Windows NT + 2000
- WWW, DNS, SMTP, POP, IMAP, Listserv, FTP, NNTP, Web proxy

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

The monitoring systems

- Mostly homegrown, open-source or customised commercial
- Link status with ping, link load with SNMP
- Alarms on web page and through email

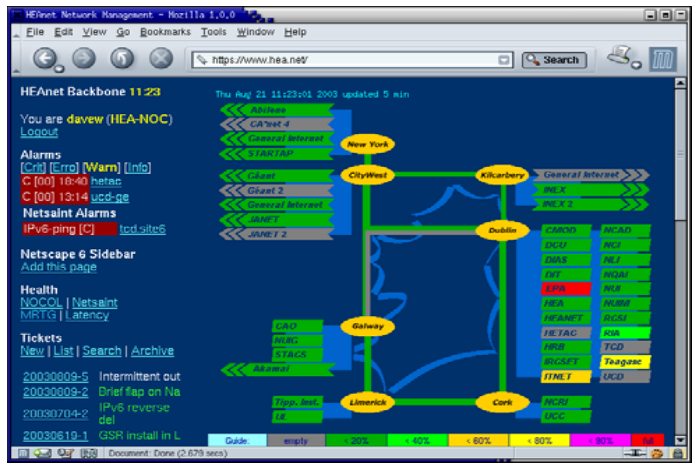

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

The people

- July 2001:
  Just moved to
  new premises
  15 staff, 10 technical
  One NOC team

- July 2003:
  Final signoff on production service
  23 staff, 18 technical
  Three technical departments

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

The plan

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

The timeline

- Step 1: Single router, single staff member, tunnels
  ➔ July–December 2001
- Step 2: Cross-city link, some NOC participation, first native links
  ➔ January–December 2002
- Step 3: Dual stack network, full staff support
  ➔ January–June 2003

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

First steps

- Needed staff
  ➔ Hired Orla McGann
- Needed IP space
  ➔ /48 in 3FFE:: from Géant GTPv6
- Needed connectivity
  ➔ Tunnel from GTPv6
- Needed infrastructure
  ➔ Assigned a 7200 + server PC

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Experimental deployment

- Setting up tunnels
- Configuring BGP4+
- Try out support on FreeBSD
- Deploy Apache 1/2, BIND 9

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

First experiences

- LIR for new address space, assigned to infrastructure and some LANs
- Connected first customer (TCD) over a tunnel
- Applied for RIPE space, set up additional tunnels to 6bone

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

First experiences

- People stole our AS number!
  (mistaken assumptions by customers)
- Started to notice the different behaviour of IPv4 and IPv6 BGP peers
  (e.g. filter lists)

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Going native

- Tunnels only take you so far
  ➔ not production-like
- Needed native links without affecting production IPv4 network
- ATM is excellent for this
  ➔ 2Mbps link across city, second router
  ➔ 2Mbps Transatlantic to 6TAP
  ➔ Used existing ATM paths

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Going native

- Now had enough experience to assess the impact on the network in the future
- Affect purchasing decisions, make feature requests (three year kit lifespan)
- Assist in dual-stacking INEX
- Beta-test for some suppliers

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Managed services

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Dual stacking services

A number of factors helped us here

- Wide spread of skills – not totally specialised
- Budget to replace old kit appeared
- Interest among staff

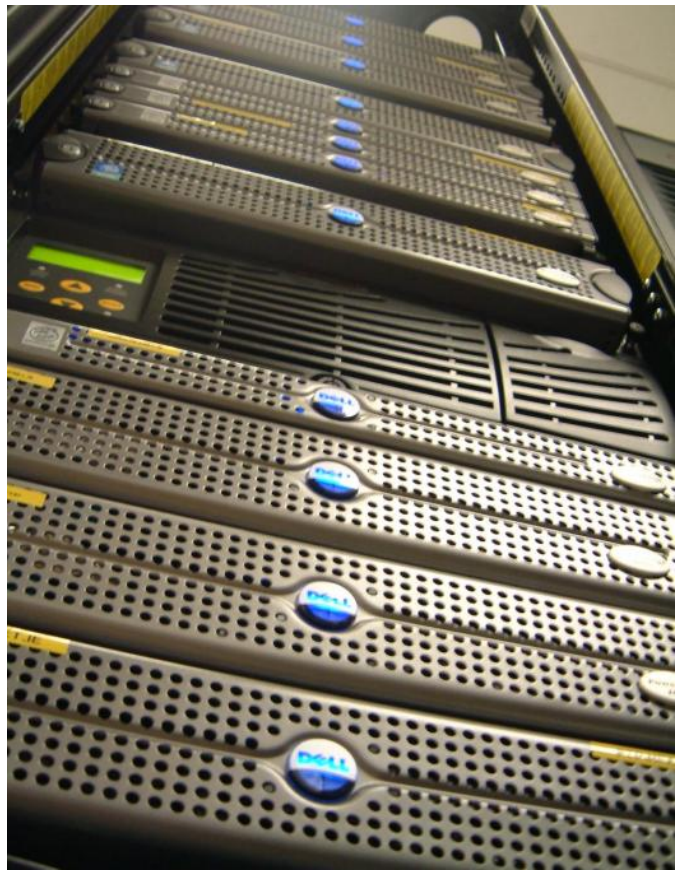

Dedicated box

- Started with a dedicated DNS & web server, zone `ipv6.heanet.ie`
- Turned on router advertisement on office & server VLANs, using separate router for IPv6
- Gained familiarity with the setup, but not production yet

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

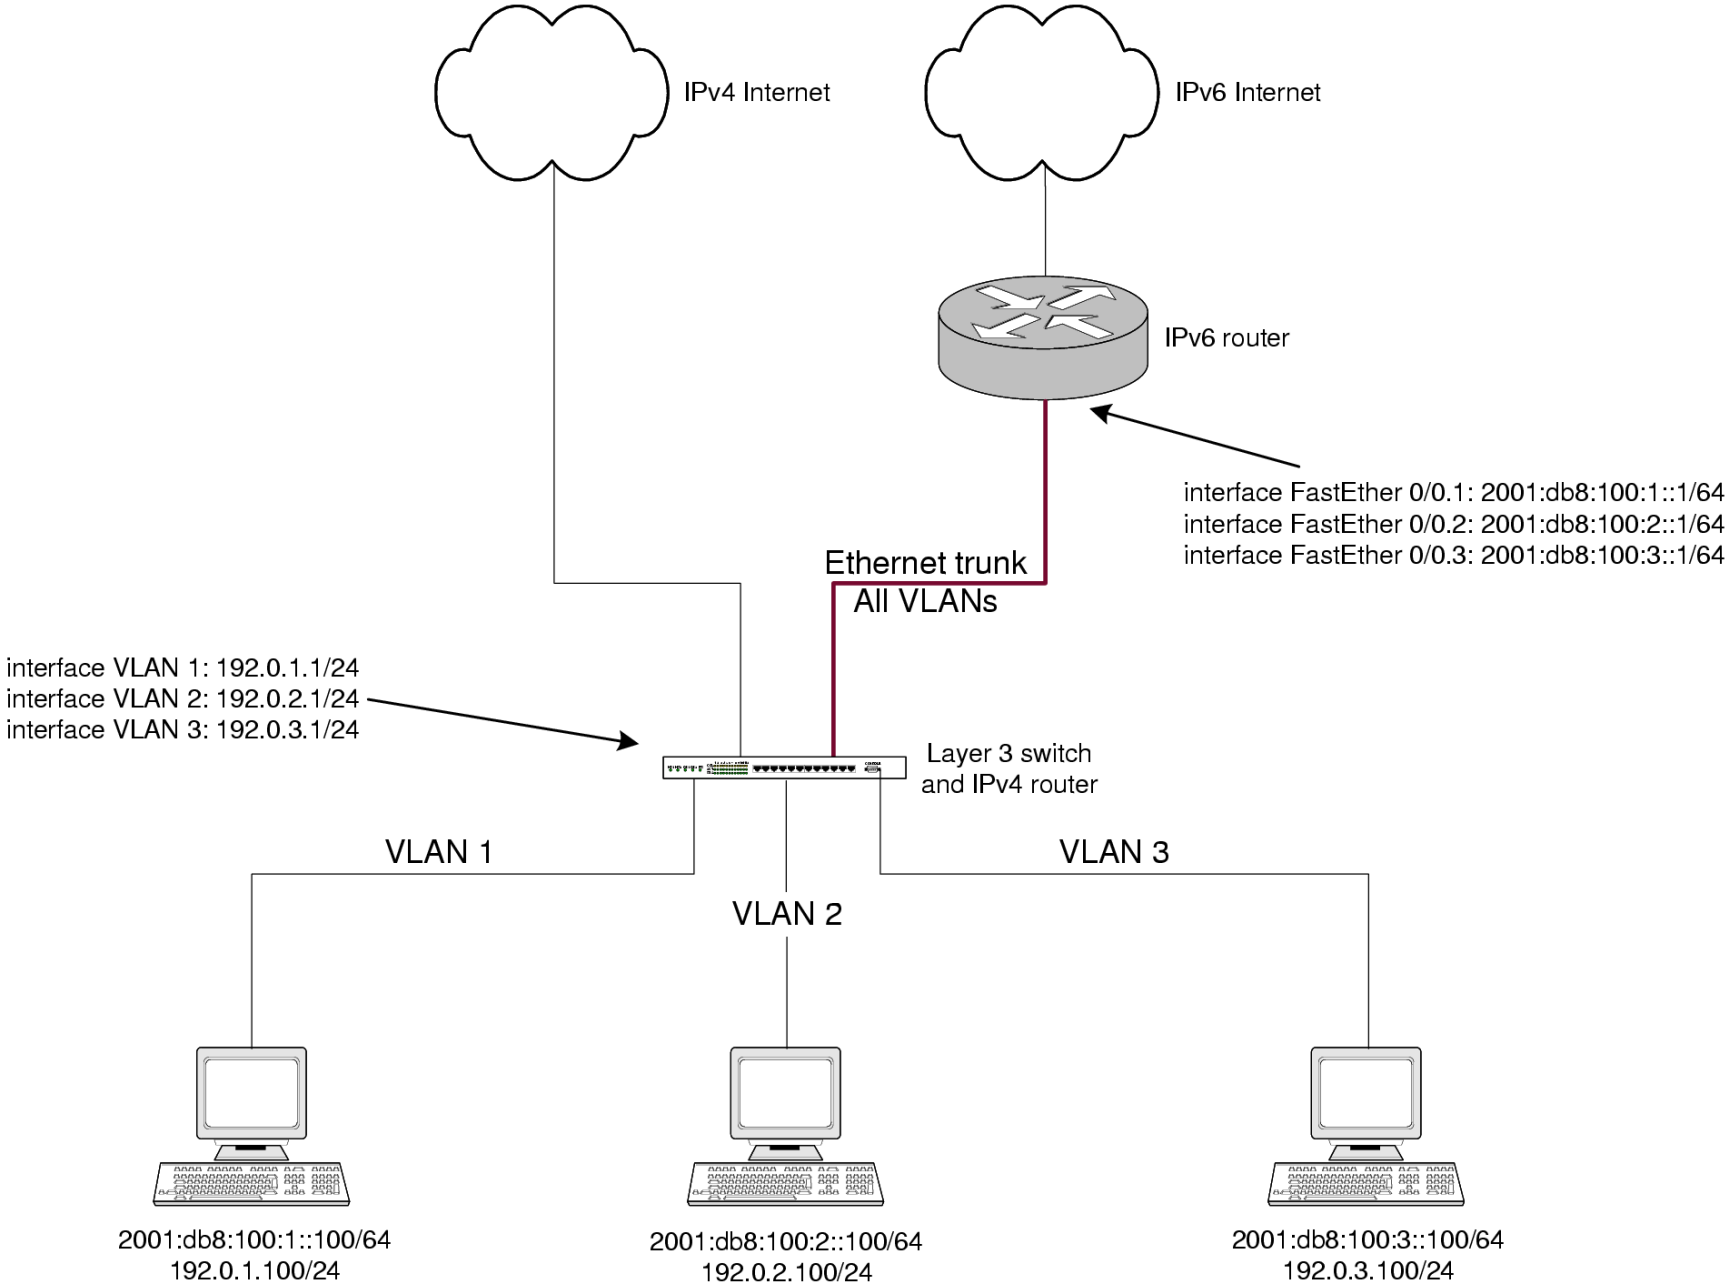
IPv4 Internet
IPv6 Internet
IPv6 router
interface FastEther 0/0.1: 2001:db8:100:1::1/64
interface FastEther 0/0.2: 2001:db8:100:2::1/64
interface FastEther 0/0.3: 2001:db8:100:3::1/64
Ethernet trunk
All VLANs
interface VLAN 1: 192.0.1.1/24
interface VLAN 2: 192.0.2.1/24
interface VLAN 3: 192.0.3.1/24
Layer 3 switch
and IPv4 router
VLAN 1
VLAN 2
VLAN 3
2001:db8:100:1::100/64
192.0.1.100/24
2001:db8:100:2::100/64
192.0.2.100/24
2001:db8:100:3::100/64
192.0.3.100/24

Router advertisement

- IPv4-only boxes are unaffected by router advertisement
- IPv6-capable boxes will receive an address and may originate IPv6 traffic
- Nothing else happens until you add AAAA records to the DNS

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Upgrading services

Classes of problems:

- OS doesn't have the stack
- Software doesn't support the protocol
- Client transition funnies
- Supporting software incompatible (e.g. log analysis tools)

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Deployment mechanism

- Migrate to new machine/install IPv6 in OS
  ➔ Immediately see IPv6 outbound
- Monitor for a bit
- Add AAAA record to DNS
- Monitor some more (depending on confidence)
- Call it production

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Daemon changes

Not all of these changes were just for IPv6

- SMTP: Sendmail → Exim
- IMAP: UW → Courier-SSL
- HTTP: Apache → Apache 2
- NNTP: innd → Diablo
- DNS: bind8 → bind9, nsd
- Listserv: depends on MTA
- Proxy: Squid → None

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Rollout

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Dual-stack the backbone

- Heterogenous environment

- Cisco 7200/7200VXRs
  ➔12.2T, 12.2S, 12.3

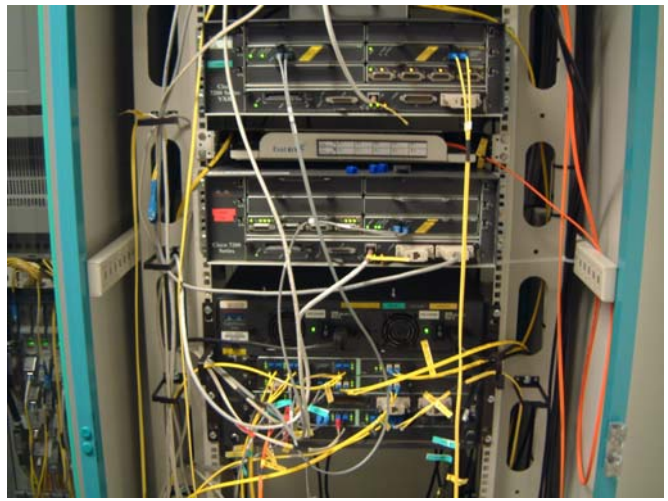

  - ➔ Tested in IPv6-only environment
  - ➔ Reasonably certain of code stability
  - ➔ Change isn't too serious

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Dual-stack the backbone

- Cisco 12000/12400
  - ➔ 12.0ST ➔ 12.0S
  - ➔ No hardware forwarding (engine 3 linecards only)
  - ➔ Consistent IOS preferred
  - ➔ Bug count, change is bad

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Dual-stack the backbone

- Cisco 7600
  - ➔ No IPv6 support (at the time)
  - ➔ Software support coming out now
  - ➔ Hardware support end of 2003 – new supervisor card

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Dual-stack the backbone

- Juniper M20
  - ➔ Full support when we needed it

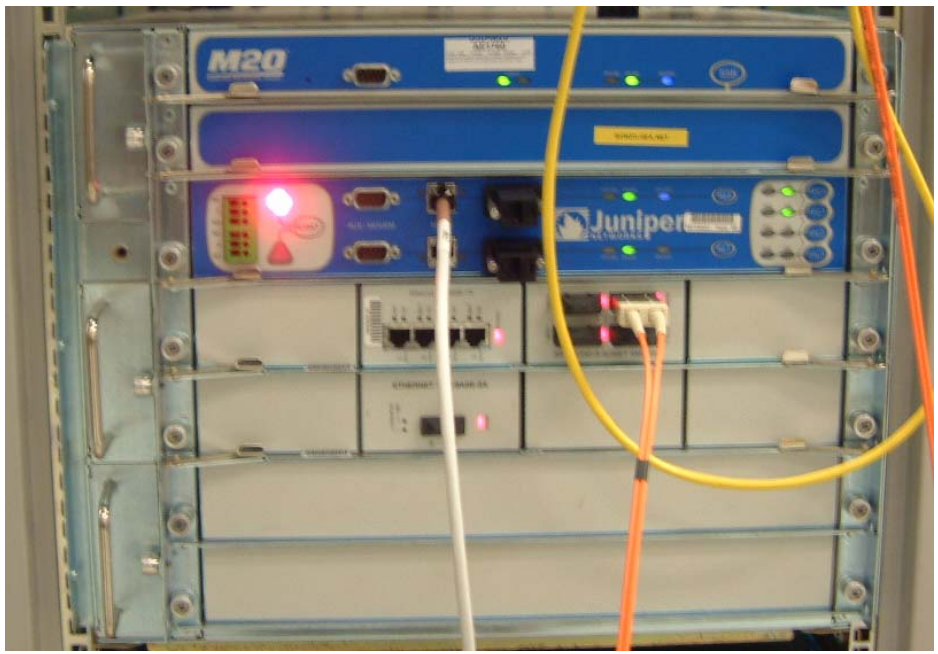

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

First production 12000

- 17 Dec 2002 – Phobos
  - Enable IPv6 routing
  - added to IS-IS and BGP mesh
  - brought up IPv6 to external peer (Abilene)
  - brought up IPv6 BGP to Abilene
  - tested routes
  - performed by IPv6 team

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

First suprise

IPv4 route maps went “missing”

```
neighbor 62.40.103.241 remote-as 20965
neighbor 62.40.103.241 description Geant STM-16 Secondary
neighbor 62.40.103.241 password 7 <removed>
neighbor 62.40.103.241 send-community
neighbor 62.40.103.241 route-map geantsec-in in
neighbor 62.40.103.241 route-map geantsec-out out
neighbor 62.40.103.241 filter-list 40 out

neighbor 146.97.40.133 remote-as 786
neighbor 146.97.40.133 description JANET
neighbor 146.97.40.133 route-map janet-in in
neighbor 146.97.40.133 filter-list 41 out
```

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

First suprise

IPv4 route maps went “missing”

```
neighbor 62.40.103.241 remote-as 20965
neighbor 62.40.103.241 description Geant STM-16 Secondary
neighbor 62.40.103.241 password 7 <removed>
neighbor 146.97.40.133 remote-as 786
neighbor 146.97.40.133 description JANET
```

...traffic unaffected

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

First suprise

IPv4 route maps went “missing”

```
neighbor 62.40.103.241 remote-as 20965
neighbor 62.40.103.241 description Geant STM-16 Secondary
neighbor 62.40.103.241 password 7 <removed>
neighbor 146.97.40.133 remote-as 786
neighbor 146.97.40.133 description JANET

address-family ipv4
neighbor 62.40.103.241 send-community
neighbor 62.40.103.241 route-map geantsec-in in
neighbor 62.40.103.241 route-map geantsec-out out
neighbor 62.40.103.241 filter-list 40 out
neighbor 146.97.40.133 route-map janet-in in
neighbor 146.97.40.133 filter-list 41 out
exit-address-family
```

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

BGP config

- Run separate v4 and v6 BGP sessions for everyone's sanity
- Activate IPv4 peers in IPv4 addr-family and deactivate IPv6 peers (+ vice versa)

```
address-family ipv4
neighbor 62.40.103.241 activate
no neighbor 2001:798:2019:10AA::5 activate
exit-address-family
```

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Choice of IGP

- Options: RIPv6, OSPFv3, Integrated IS-IS

- On Cisco:
  - ➔ OSPF runs as a separate process
  - ➔ IS-IS uses single process, and address-families (like BGP)

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Choosing an IGP

- Different IGPs for IPv4 and IPv6:
  - ➔ Separate control, separate networks
  - ➔ Separate failure modes
  - ➔ Cross-protocol problems less likely
- Same IGPs for IPv4 and IPv6
  - ➔ May get cross-pollination
  - ➔ May require same layout for v4 + v6
  - ➔ Easier troubleshooting

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Developing documentation

- Dual-stack first router
  - ➔ by ipv6 team, write up procedures
- Dual-stack second router
  - ➔ by ipv6 team, revise procedures
- Dual-stack subsequent routers
  - ➔ by operations, iterate till procs ready

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Implementation

- 17 Dec 2002 – Phobos
  - ➔ by netdev, write up procedures
- 26 Mar 2003 – Deimos
  - ➔ by netdev, minor revision
- 2 Apr 2003 – Charon
  - ➔ by netops, netdev supervise
  - ➔ continued every 1-2 weeks
- procedures handed over

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Workarounds

- GE and ATM customers land on Cisco 7609
- No IPv6 support
- Fake GE support using VLANs
- Provision separate ATM PVC

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

VLAN meddling

On the layer 3 switch:

```
interface GigabitEthernet3/3
 description TCD [Vlan101] (COLT DUB/DUB/LE-001892)
 no ip address
 no logging event link-status
 speed nonegotiate
 switchport
 switchport access vlan 101
 switchport mode access
!
interface Vlan101
 description TCD link VLAN (Physical Gig3/3)
 ip address 193.1.196.149 255.255.255.252
 no logging event link-status
!
```

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

VLAN meddling

On the IPv6 router:

```
!
interface FastEthernet0/0.101
 description Link to TCD site router
 encapsulation dot1Q 101
 ipv6 address 2001:770:8:3::1/64
 ipv6 router isis backbone
!
```

This is transparent to the customer

(but not gigabit in this case)

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

ATM workarounds

Separate PVCs over virtual path

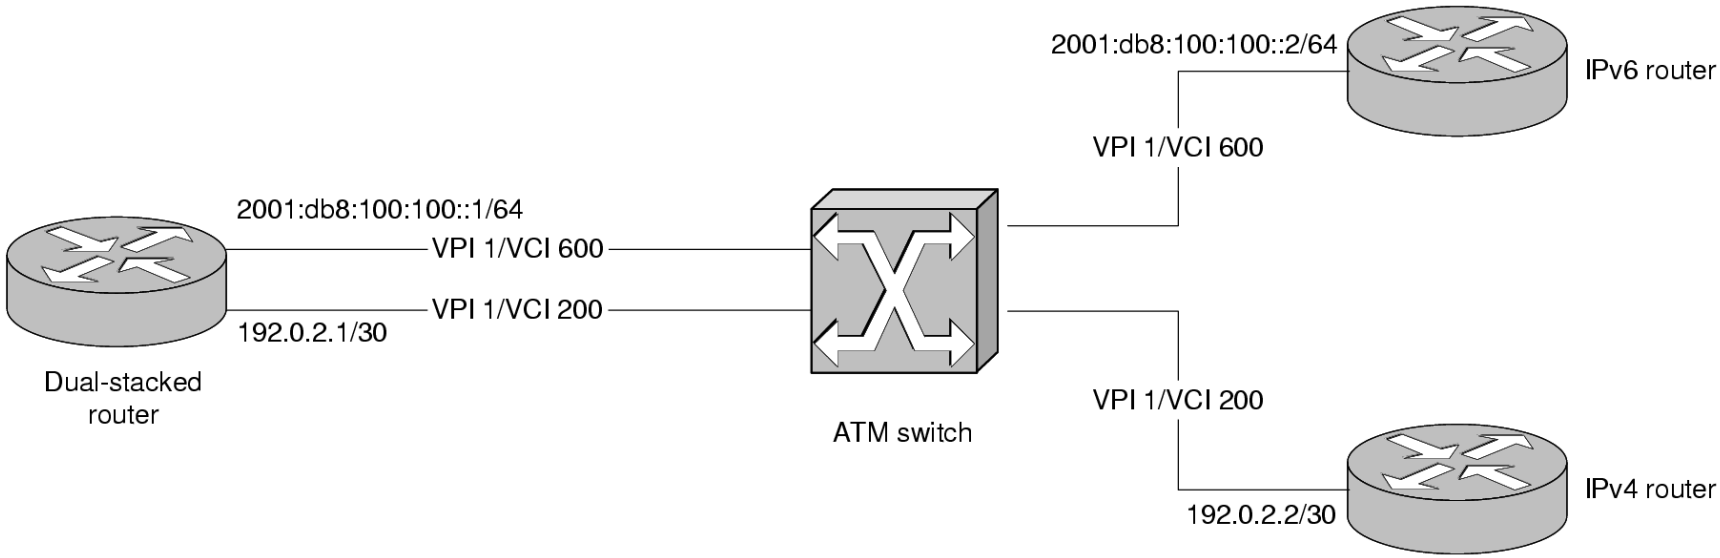

- Requires you to dedicate bandwidth to the IPv6 PVC

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Hardware compatibility

- 7200/7200VXR use software forwarding
- 7600 is dependent on Supervisor (Fourth quarter, 2003)
- 12000/12400 is dependent on linecard
  - ➔ IPv6 in Engine 3 linecards ONLY
  - ➔ Other linecards use software fwding
  - ➔ Depends on ingress card

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Management

- SNMP over IPv6 often not supported yet
  ...but this doesn't stop you monitoring
- On Cisco and Juniper, interface counters show total traffic, not per-protocol
- VLAN + ATM workarounds can separate this traffic out

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Routing funnies

- Tunnels are *really really* annoying
  - ➔ Maximum throughput limit
  - ➔ Extra latency with every hop
  - ➔ Routing policy not in RIPE DB, planning is difficult

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Routing funnies

- Instituted more restrictive tunnel policy:
  - ➔ Direct peers, yes
  - ➔ Networks with no IPv6 access, yes
  - ➔ Otherwise no
- RIPE Test Traffic IPv6 project
- Looking glasses, maps

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Routing funnies

- Your IPv6 routing policy might not match your IPv4 routing policy

```
*>i2001:610::/32     3549 1103 i      GLBX
* i                  3425 1103 i      6TAP
* i                  20965 1103 i     Géant
* i                  11537 1103 i     Abilene
* i                  10764 1103 i     STAR TAP
```

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

External connectivity

- Transit in US from Abilene
- Transit in EU from Géant and Global Crossing (tunnel)
- Peering with local ISPs at INEX
- Roughly matches our physical & IPv4 layout

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

External connectivity

- Little control over remote path

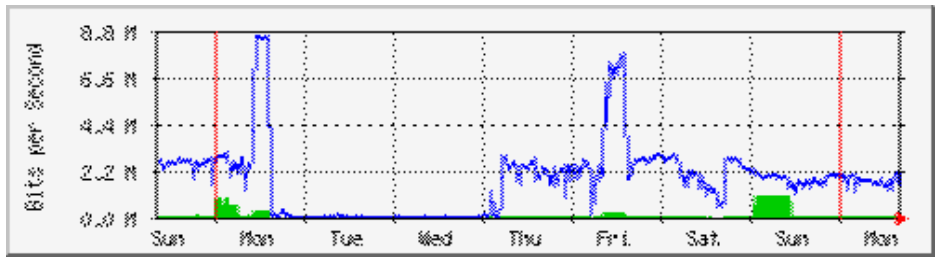

- If you don't have direct connection, the intervening hops might not be up to it

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Routing uglies

- Ghost routes for old /35s
- Visible as absurdly long AS paths
  - ➔ Gert's talk to ipv6-wg
- We haven't taken this leap yet

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Observations

- Typical traffic – small but *very* bursty
- Customers suddenly asking for connectivity with no notice

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Approaching signoff

- Completed internal documentation
- Tutorial and test lab for staff

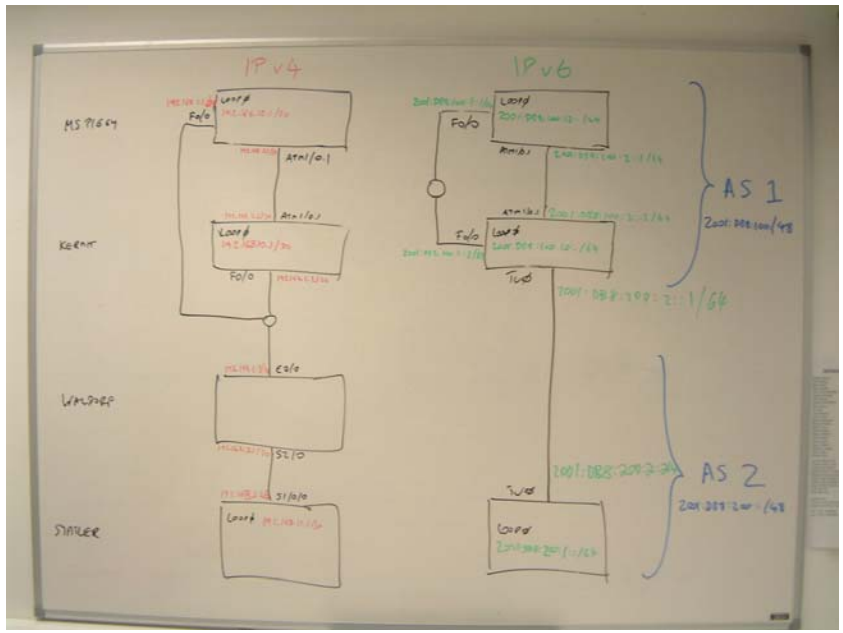

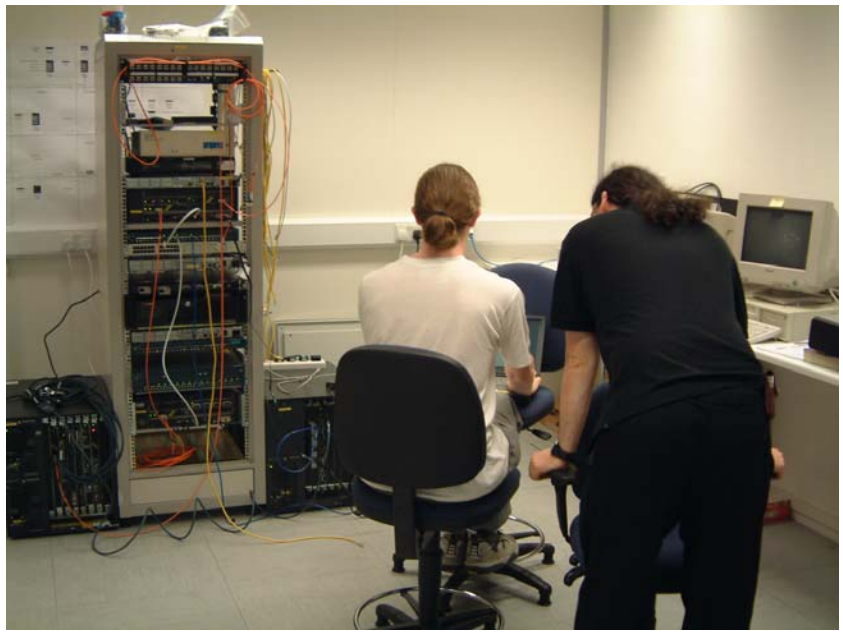

IPv6 LAB.

1. BRING UP TUNNEL BETWEEN KERMIT + STATLER
   ADDRESS RANGE 2001:DB8:200:2::/64, SOURCE LOOPBACK INTS.

2. BRING UP IS-IS BETWEEN KERMIT AND MSPIGGY
   INCLUDING LOOPBACKS

3. BRING UP IBGP (AS 1) BETWEEN KERMIT + MSPIGGY

4. BRING UP EBGP (AS 2) TO STATLER

5. PING 2001:DB8:200:1:: FROM MSPIGGY

Production and Signoff

- Final router dual-stacked 16 July 2003
- Cisco DoS announced 17 July 2003
- Final signoff 29 July 2003
- 2 native customers, 2 tunnelled, more delegated/asking

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Lessons

- Took time but not really money
  - ➔ but purchasing decisions now will affect performance in 2006
- Customers do not give advance notice of demand
- Some stimulation of interest nationally

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Questions?

PLEAN FORBARTHA NÁISIÚNTA

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK