APNIC Operations
IPv6 Challenges

Challenges

- Network/System structure/inheritance
- Management desires for adoption of IPv6
- External parties, provisioning

Let's talk about the Operations area

Technical Challenges

- Network
  - IP addressing
  - What is the V6TE?
  - Topology
  - Connectivity/Upstream availability

Technical Challenges (cont)

- Systems
  - IPv6-ready applications
- Monitoring System
  - What's possible and what's not?

Our Network

- IPv6 addressing
- What is the V6TE?
- Topology
- Connectivity / upstream availability

IPv6 Addressing

- We have a single IPv6 /32 sub-divided into a /35 per site
  - Five sites + 'infrastructure' + a special-project, subnet (V6TE)
  - This required the creation of individual whois route objects for each subnet:
    - An inet6num and a route6 per /35

```
inet6num:    2001:0DC0:4000::/35
netname:     APNIC-AP-V6-HK
descr:       APNIC Pty Ltd - HongKong
network
descr:       Level 1, 33 Park Rd
descr:       PO Box 2131, Milton
descr:       Brisbane, QLD.
country:     AU
```

```
route6:      2001:dc0:4000::/35
descr:       APNIC-IPv6 HONG KONG
origin:      AS18366
mnt-by:      MAINT-APNIC-AP
changed:     hm-changed@apnic.net 20090226
source:      APNIC
```

IPv6 Addressing (cont)

- We inherited a legacy subnet model
  - IS there really a ‘good’ one?
- What we did was model from /35 to /48, to /64
  - Addressing combination = /64 using the last 3 HEX digits = value from the IPv4 ‘dotted quad’ address
  - This makes it easy to configure the servers and DNS entry
- Example
  - IPv6 Address : 202.12.29.197 (icons.apnic.net)
  - IPv6 Address: 2001:dc0:2001:11::197/64

About our Tunnel Broker (V6TE)

- Why we set up the V6TE (IPv6 Transit Exchange)
  - The purpose was to help Asia Pacific economies/account holders gain experience with IPv6, by giving them access to the IPv6 cloud via a tunnelled protocol
- But now...
  - We have more in-region native IPv6 providers
  - Sub-optimal routing maintained by V6TE tunnels

V6TE History

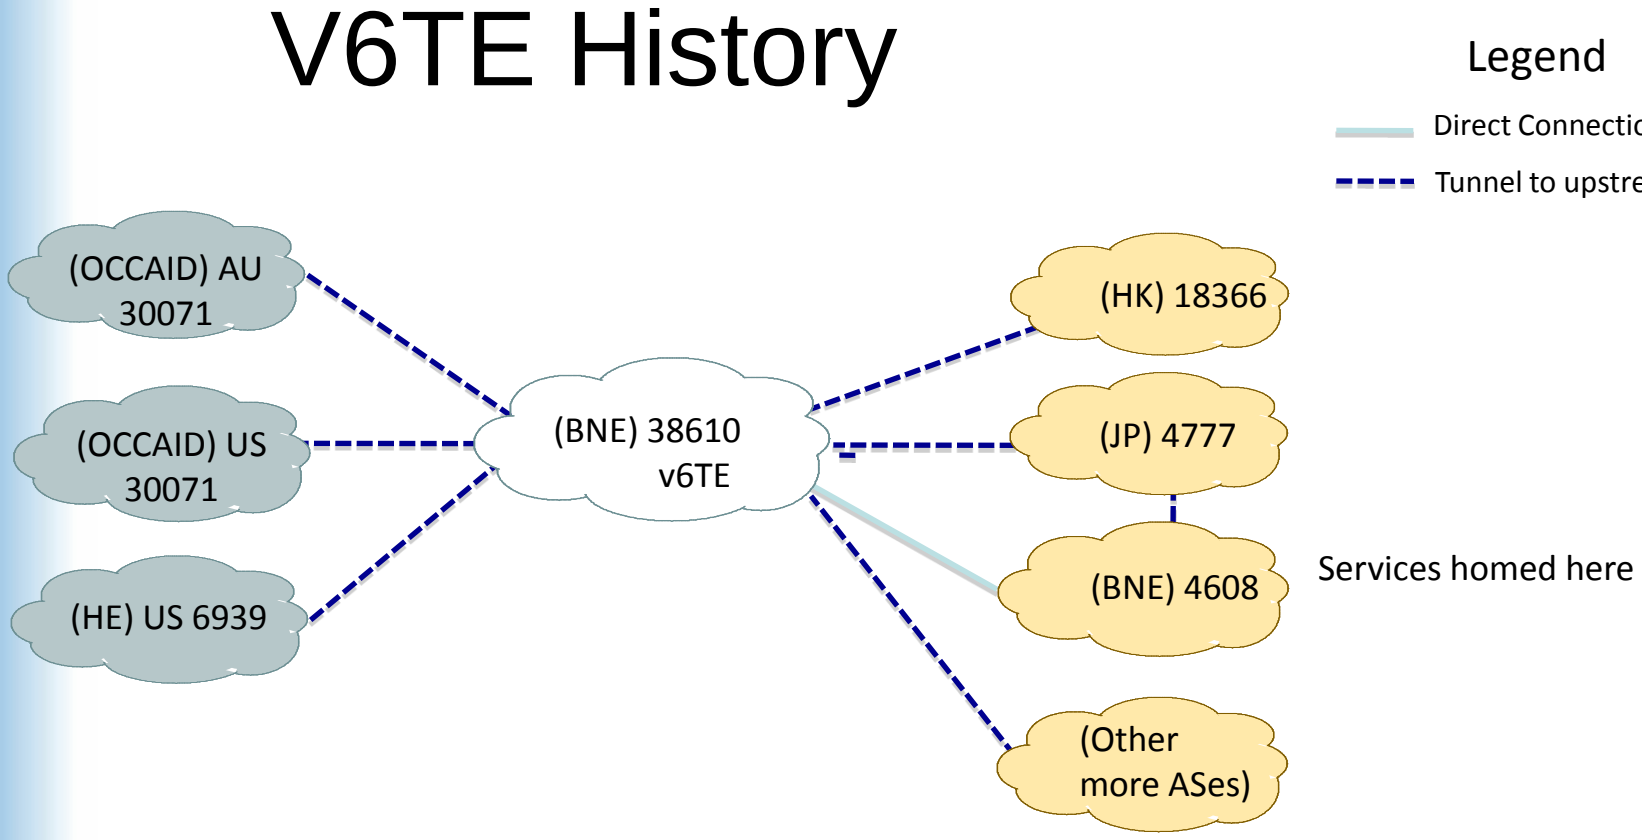

- Our IPv6 network in Brisbane was routed via this project subnet, using a tunnel
- Since native IPv6 hit local IX, we have moved APNIC service routing off V6TE
- We are considering the long-term future of this service

Current V6TE Users

- New Zealand = 10
- Australia = 3

Native IPv6 now available at some IX

- Philippines = 3
- Pakistan = 2
- China = 2
- Singapore = 1
- India = 1
- Thailand = 1
- Vietnam = 1

With HE offering a tunnel broker service it may be wise for these participants to peer with them directly instead.

Topology

- This is one of our major challenges, since none of our externally visible sites are directly connected.
  - (This is for our JP, HK & US located nodes)

Topology (cont)

- In Brisbane we (currently) have two sites:
  - A back-office and co-location (co-lo) site, both of which connect to upstreams, and have dark-fibre (private) connections
  - These shared the same /35 initially, which required us to renumber using a separate block of /35 for each.
  - No v6-aware IMS, so ‘IPv6 renumbering is easy’ is not quite true:
    - Lots of DNS, route, and configurations to change by hand.

Old Structure

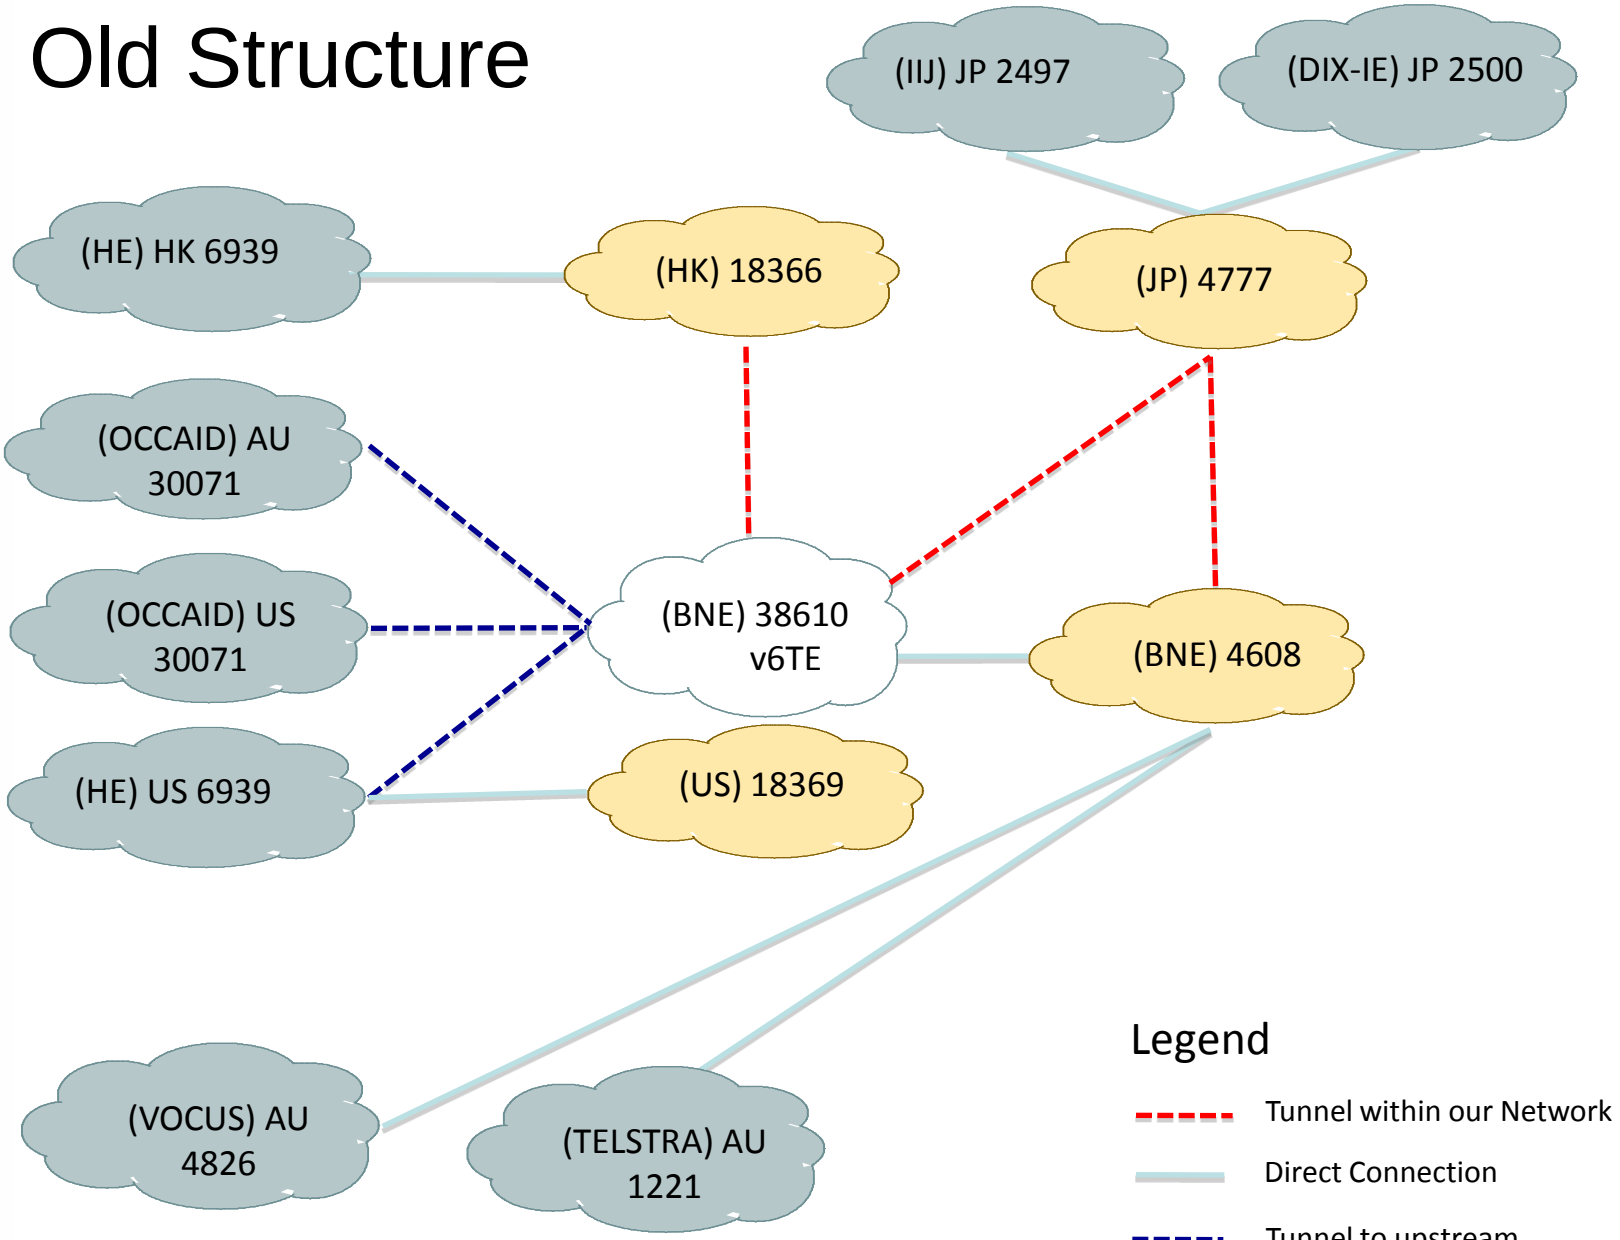

Connectivity (JP/HK/US)

- Japan - no problems getting IPv6 access
  - NSPIXP6 was our first native IPv6 association
  - IPv6 peering is easily available in KDDI/NTT Otemachi rooms
- Hong Kong - also no problem
  - HKIX have IPv6, and HE is there
- United States - we had some problems initially
  - Our upstream required us to join a research project related to IPv6 before they would give us a connection
  - We could find no further details about this project.

Connectivity (Brisbane)

- Historically difficult to get native IPv6 connectivity
  - Until Vocus came into the picture last year
- Other problems
  - Telstra re-announcing the same prefix to us from our v6TE
  - Soul/TPG don't have IPv6 connectivity to give us
- We managed to survive with tunnels for a number of years
  - Thanks to HE and OCCAID

Connectivity (Brisbane)

- Brisbane network last to have native access compared to our other sites
  - Therefore service delivery via IPv6 delayed, except DNS (at other nodes)
  - But is our main investment in co-location for service delivery of mail, web, ftp, dns
- Thanks to:
  - HE
  - Vocus
    - For improvements in IPv6 availability; access to our Brisbane node

Our IPv6 Topology (Full mesh using some tunnels)

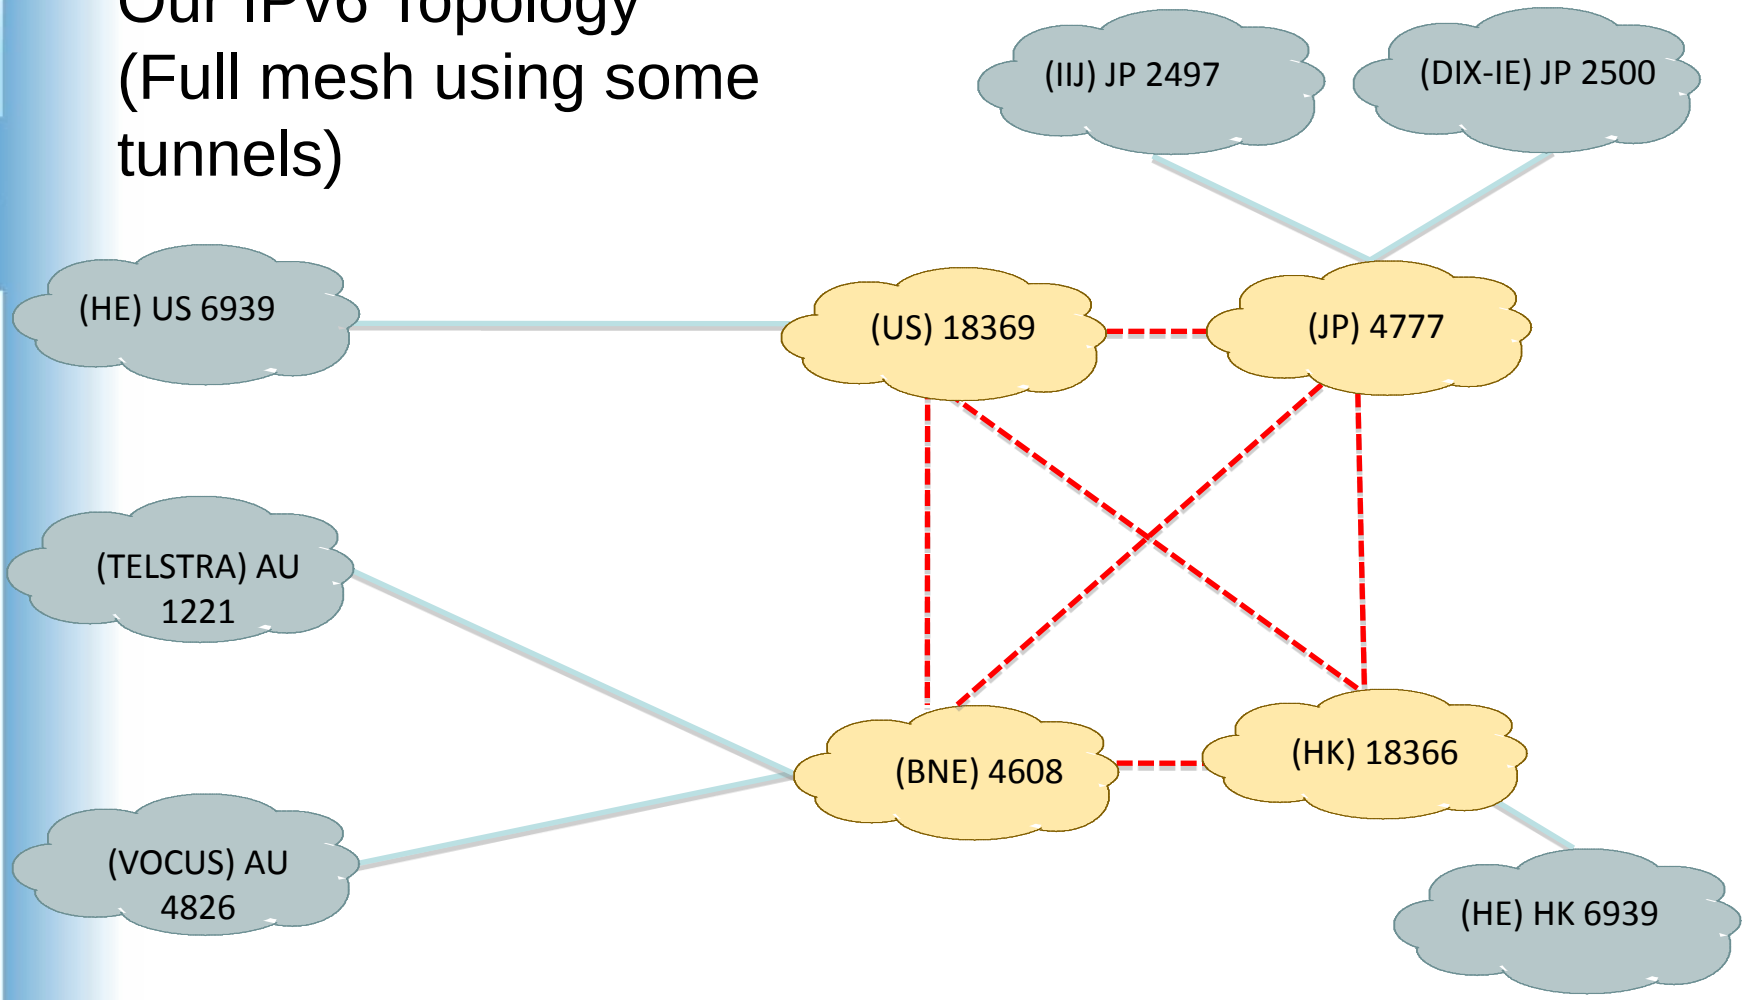

Why do we still have tunnels? (see the next slide)

Legend

- Tunnel within our Network
- Direct Connection
- Tunnel to upstream

Why do we still have tunnels?

- We don't have direct connections to those key sites in JP, HK, and US
- Why not use the IPv6 cloud?
  - The tunnel only serves routes for the /35s per site.
  - We have good evidence the tunnels improve our own transit exchange
- Brisbane (4608) announces the /32 aggregate
  - Including the individual /35 to Vocus
  - We had problems initially since they do not allow 35's
    - If it's in an IRR, shouldn't it be permitted?

Sample traceroute6 from AU to HK

- With tunnel: **150ms** RTT, consistent.
- traceroute6 2001:DC0:4001:1:0:1836:0:1
- traceroute6 to 2001:DC0:4001:1:0:1836:0:1 (2001:dc0:4001:1::1836:0:1) from 2001:dc0:a000:4:21e:c2ff:fe06:f94c, 64 hops max, 12 byte packets
- 1. 2001:dc0:a000:4::1 0.394 ms 0.555 ms 0.344 ms
- 2. 2001:dc0:e002:4608::2 0.849 ms 0.639 ms 0.662 ms
- 3. 2001:dc0:4001:1::1836:0:1 148.994 ms 148.958 ms 149.085 ms

Sample traceroute6 from AU to HK (cont)

- Without the tunnel: **300ms** RTT, variable routing.

traceroute6 2001:DC0:4001:1:0:1836:0:1

- traceroute6 to 2001:DC0:4001:1:0:1836:0:1 (2001:dc0:4001:1::1836:0:1) from 2001:dc0:a000:4:21e:c2ff:fe06:f94c, 64 hops max, 12 byte packets
- 1. 2001:dc0:a000:4::1 0.578 ms 0.330 ms 0.413 ms
- 2. 2001:dc0:e002:4608::2 0.797 ms 0.664 ms 0.643 ms
- 3. as4826.ipv6.brisbane.pipenetworks.com 1.046 ms 0.922 ms 0.791 ms
- 4. ge-0-0-1.bdr01.bne02.qld.vocus.net.au 0.925 ms 1.229 ms 1.108 ms
- 5. ge-1-0-5.cor01.syd03.nsw.vocus.net.au 18.418 ms 18.549 ms 18.811 ms
- 6. ge-0-0-0.bdr01.syd03.nsw.vocus.net.au 18.899 ms 18.811 ms 18.731 ms
- 7. 2001:de8:6::1:26:1 17.541 ms 16.326 ms 15.788 ms
- 8. 2001:cb0:a102:1::2 141.088 ms 141.073 ms 140.872 ms
- 9. 2001:cc8:102:17::1 141.338 ms 140.906 ms 141.098 ms
- 10. 10gigabitethernet2-3.core1.sjc2.he.net 218.593 ms 218.603 ms 216.704 ms
- 11. 10g-3-2.core1.sjc1.ipv6.he.net 218.068 ms 221.591 ms 224.625 ms
- 12. v1026.core1.hkg1.he.net 404.758 ms 409.540 ms 407.768 ms
- 13. * * *
- 14. * * *
- 15. 2001:dc0:4001:1::1836:0:1 304.566 ms 303.111 ms 303.461 ms

Systems (Services)

- Some applications were not ready
  - We had issues with our MyAPNIC service
    - Apache ‘session ID’ length derived from src IP
    - Database field was too small (32); we needed to increase to accommodate IPv6 src
- Issues with our www services running IPv6
  - Reverse-proxies (HA) via Squid running on Redhat
    - IPv6 support was by patch only, not a supported product (HA compromise)
  - Load balancer IPv6 support (RADWARE)
    - Initially was pre-release test image only

Systems (Services) (cont)

- Whois service
  - RIPE whois codebase in use did not have native IPv6
    - Deployed a server that responds to IPv6 queries (jwhois.apnic,net)
    - However, its responses are not identical to core whois
    - Format of responses differ by IP protocol.
    - New version of RIPE whois with native IPv4 and IPv6 is in test phase

Systems (Services) (cont)

- Reverse DNS service
  - DNS servers work well, but subject to initial routing problems; instabilities in host (Redhat Linux) IPv6 stack
  - Delegation engine works well with IPv6, some confusion about delegation boundaries for IPv6 reverse with clients
    - Need to improve documentation on the semantics of ‘dots’ embedded in zone files, the delegation boundaries we support

Systems (Services) (cont)

- Helpdesk (web-based chat)
  - Works without a problem. Cleanly abstracted from IP transport
- APNIC meeting support
  - Jabber service for the conference required platform upgrade
  - Video streaming via web only for IPv6
  - Quality of ad-hoc (conference) v6 routing variable

Systems (Services) (cont)

- We host the 6to4.NRO.net web server
  - Reverse delegation request for 6to4 address blocks
    - Service runs best if bound to a 2002::/16 prefix (ensures clients detected from origin 2002::/16 src IP)
    - We had problems here because 6to4 v6 prefix was not properly routed inside our own network
    - Being ‘multi-homed’ on different IPv6 complicates things

Systems (Services) (cont)

- Email system works
  - With some issues on access-list since we are doing gray listing by IP as well as from “address”
- Internal web services ok
  - Choice of v4/v6 transport mostly hidden to users

System (Services) (cont)

- Network Share (SAMBA)
  - We had problems with our network share because of the DNS query acquiring an IPv6 address
  - We are doing combination EUI-64 and DHCPv6 on the router
    - Breaks IPv4 access to the share if IPv6 is enabled
      - (on MSWindows hosts only)

Monitoring

- Monitoring of our services with IPv4 is easy
  - Well understood, scripted tests available
- But with IPv6 there is a challenge
  - We are using an Open Source monitoring system (Zenoss Community Edition)
  - Currently monitoring all our critical services
    - By default, no support for IPv6, but we have a work around

Zenoss with IPv6

- It has the mibs, but requires customization as
  - It's not enabled by default
    - IPV6-ICMP-MIB.txt
    - IPV6-MIB.txt
    - IPV6-TC.txt
    - IPV6-TCP-MIB.txt
    - IPV6-UDP-MIB.txt
  - Once configured, basic IPv6 monitoring worked
- Using customized scripts taken from Nagios
  - (Backwards compatibility dependency)

Query scripts

- Nagios binaries
  - Variable query using “-6” type switch
    - » Plugin like: check_http_-6
- We are relying on DNS IPv6 records to make it work
  - Zenoss ---> DNS (to get the host details) ---> then process the script
    - Downside - if the DNS is not working the IPv6 monitoring won’t work
- This is only a work around until Zenoss fully supports IPv6

Summary

- Avoid testbed/pilot/legacy confusion
  - Think about your address plan, reverse DNS early in deployment, or be prepared to re-deploy from scratch
- Tunnels are bad… but also useful
  - Cleanly separate in routing architecture
  - Avoid re-announcing tunnels as higher-pref routes over native alternatives
- Application test in depth
  - Simple checklists like “is webserver v6 enabled?” won’t reveal problems

QUESTIONS?