



# IPv6 Network Operators BOF

Stewart Bamford

*Stewart.Bamford@level3.com*

NANOG 39 – Toronto

February 4<sup>th</sup> to 7<sup>th</sup> 2007

# Introduction & House Rules

- This is a BOF for people who:
  - Run an IPv6 network
  - Intend to run an IPv6 network
  - Want to discuss some “best practice” ideas
- This is **NOT**:
  - Another opportunity to argue for or against IPv6
    - » At least not until the end, then go for it, along with multihoming!

IPv6

# Aims & Approach

- The purpose of this BOF is to discuss some of the main operational issues of an IPv6 network.
- Slightly different approach to a BOF. The slides are meant to provide a framework for discussion.
- This BOF **needs your input**.
- Please feel free to interrupt at any time with questions or comments.

IPv6

# Basic Agenda

- IPv6 – How do you do it?
- IPv6 peering
- IPv6 tunnelling
- IPv6 filtering
- Supporting IPv6
- IPv6 network capacity
- Future IPv6 issues and status

IPv6

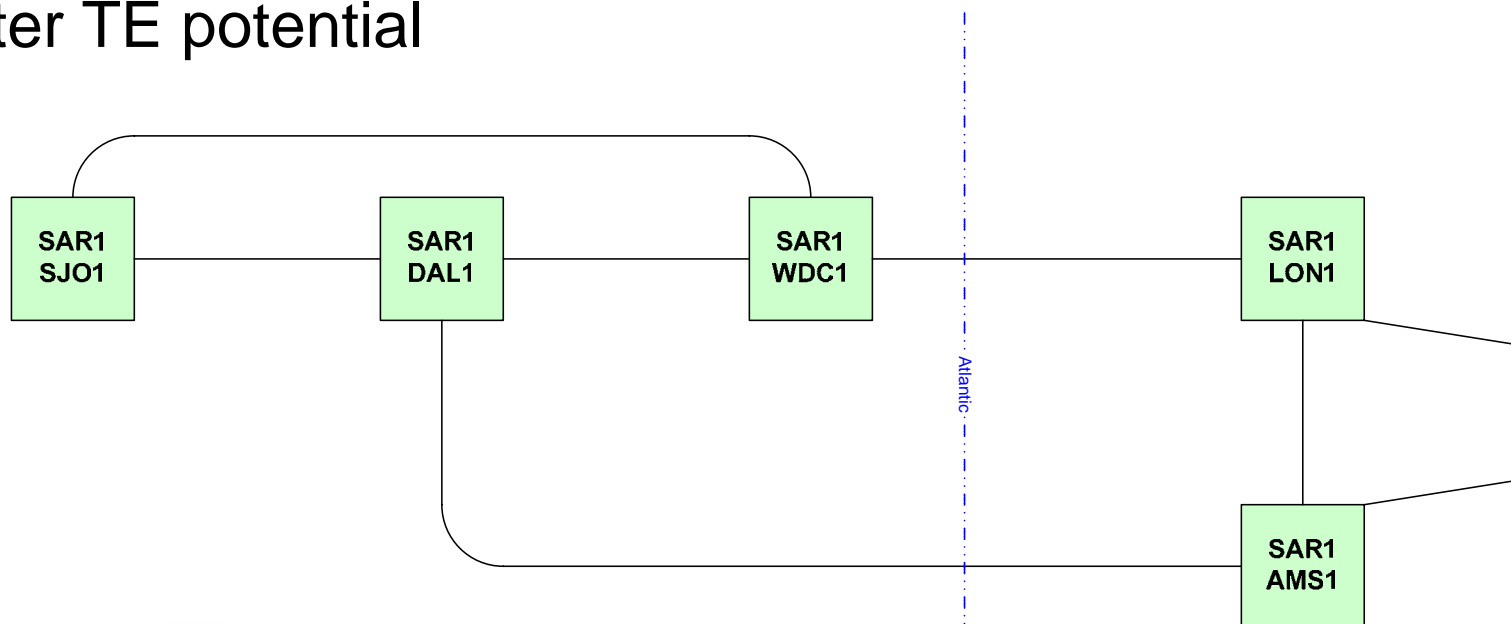
# IPv6 – How Do You Do It

- What mechanism do you use
  - Tunnels, 6PE, dual-stack?
- Why did you choose this?
- What issues did you run into?
- Would you do it again?
- Level 3 recently moved away from 6PE with expansion into more cities.

IPv6

# Level 3 IPv6 - 2007

- Now in more cities & using OSPFv3
- More peering
- Better TE potential



# IPv6

- Most “regular” BGP communities are supported in IPv6, including blackhole.
- Trying to follow the IPv4 model.

# Tunnels – A Necessary Evil?

- Problems with IPv6 over IPv4 tunnels:
  - Vendor issues
  - Scaling issues (config, systems, management etc)
  - One protocol relies upon the stability of another!
- Good points about tunnels:
  - Easy to configure, at least at the start.
  - Don't require any changes to the underlying IPv4 network.

The text 'IPv6' is rendered in a large, blue, 3D-style font with a gradient and a drop shadow effect.

# Tunnels

- Do you use tunnelling?
  - Standards (best practice) configuration
  - /64, /126 or something else?
  - Long distance tunnelling
  - GRE, IPv6IP?
  - Supporting tunnels in the real world. Do they cause you problems?
    - Following slides are tunnel issues from real world situations.

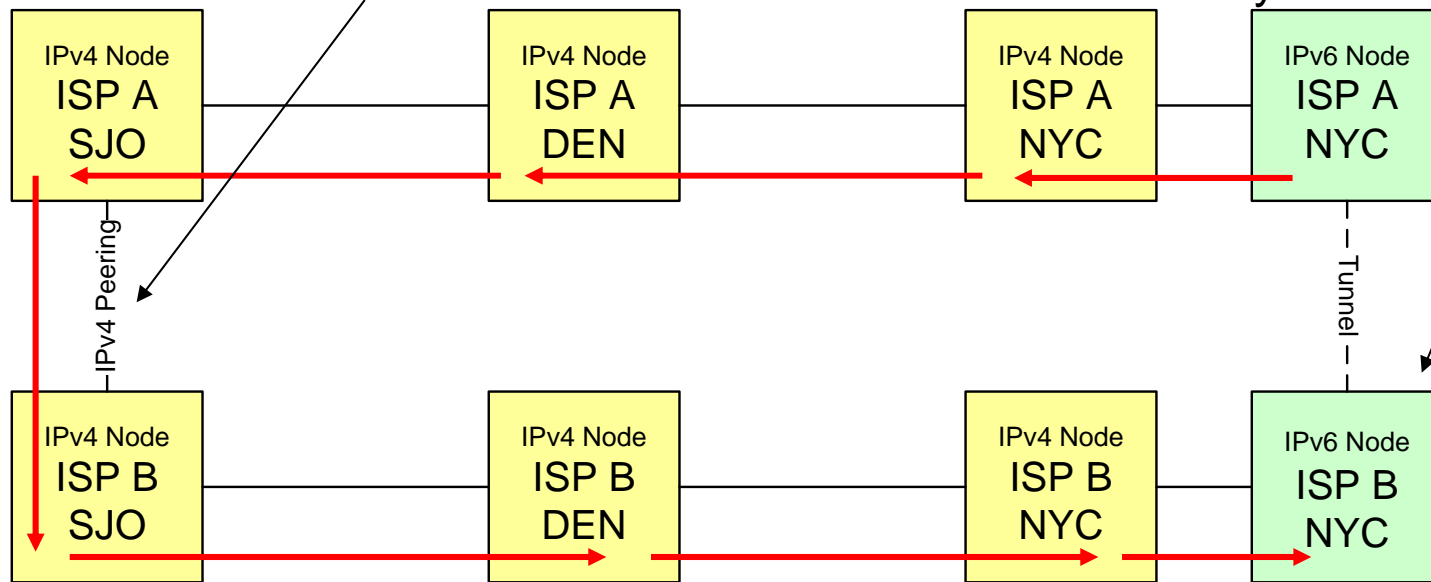
IPv6



# Tunnels – Example Issues

Only working IPv4 peering between ISP A and ISP B is 1000's of miles away

Two nodes in same building. Could be <1ms latency.



# IPv6

- Geographical location is **not** an indication of tunnel path. Actual path and latency of tunnel can vary depending upon stability of IPv4 network.



# IPv6 Peering

- Is it just the same as IPv4 peering?
- Do the same rules and policies apply?
- Public or private?
  - Where's best (for new operators)?
- Native, dual stack, tunnelled?
- Transit via peering. A good thing?

IPv6

# Filtering

- Some IETF, RIPE etc stuff on this before.
- Doesn't seem to be too detailed.
- No real agreement.
- Should we even bother agreeing on some form of standard?
- Document last year:
  - <http://www.cymru.com/Bogons/ipv6.txt>
    - » Written by KDDI & co
    - » Who's read this?

IPv6

# IPv6 Packet Filtering

- Should we even bother?
- Basic proposal seems to be:
  - Reject packets with source address:
    - Bogons (loopback, unspec, default, site local etc)
    - That are part of your own prefixes (where applicable)
- Usual issues with ingress packet filtering.

IPv6

# IPv6 Prefix Filtering

- Going to be necessary to do *something*
- What do folks do now?
- Level 3 approach at the moment:
  - Devices now in NA and EU. BGP and OSPFv3.
  - Don't do any special **packet** filtering.
  - **Prefix** filtering done inbound and outbound.
  - Max-prefix on customers and peers also.

IPv6

RIPE OBJECT: **AS-LEVEL3-IPV6**

NETWORK OPERATORS BOF – NANOG 39

# IPv6 Prefix Filtering (2)

- ‘Validate’ all IPv6 prefixes in and out.
  - No martian prefixes (communities or AS’s either)
  - Check prefix against IANA list (requires updating)
    - “Strict” method. Could just allow general valid IPv6 global unicast.
    - <http://www.iana.org/assignments/ipv6-unicast-address-assignments>
  - Only send our own aggregate prefixes, not specifics.
  - Don’t allow our own aggregates back in, ever.
  - Don’t allow anything from our ranges back in from peers.

IPv6

# IPv6 Prefix Filtering (3)

- Don't accept prefixes longer than.....
  - Many answers to this question
  - /32, /35, /48, /128 ??
    - ARIN seems to allocate /48 directly to end-users (6.5.8.2)
  - Inherently linked to the multi-homing “issue”
  - Level3 approach is to allow in up to a /48
    - Yes, that allows some traditional multi-homing!
    - Don't leak anything out of our aggregates out though
    - You mess with your space if you like, we'll try and keep ours tidy 😊

IPv6



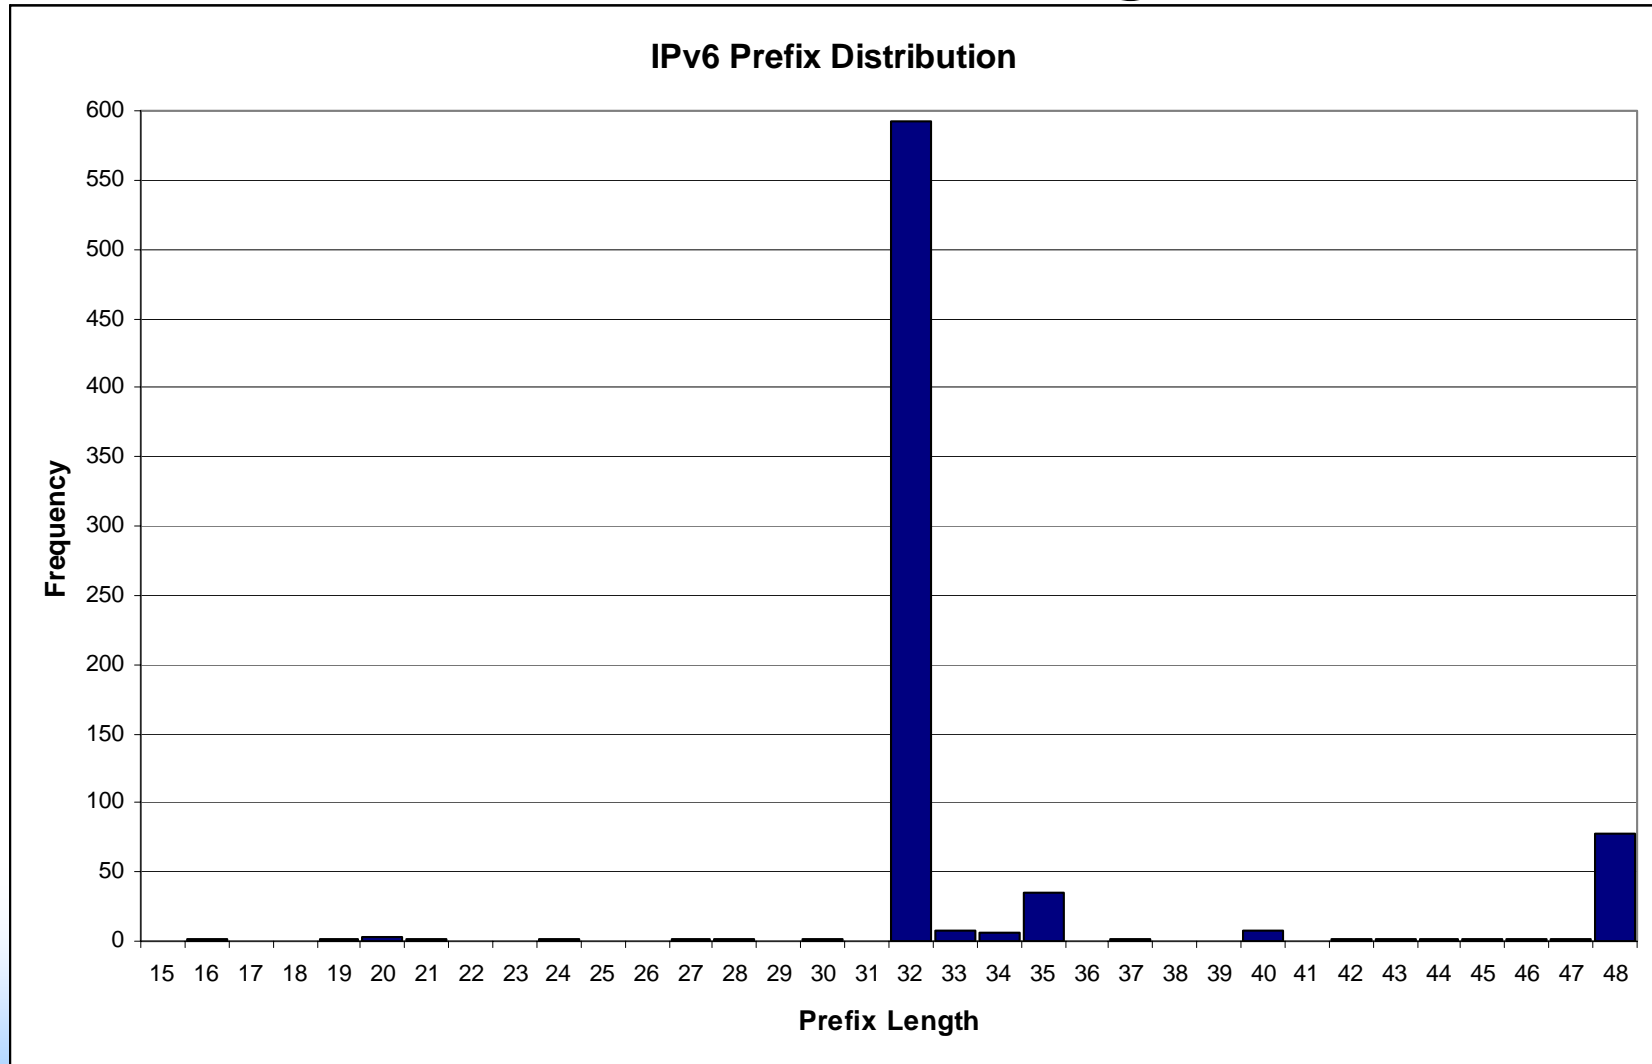
# IPv6 Prefix Lengths

- Data from live IPv6 table.
- Taken from 'sane' prefixes.
- Shows a majority of /32.
- Spike of /48 around.

Length	Freq	%age
16	1	0.13%
19	1	0.13%
20	3	0.40%
21	2	0.27%
24	1	0.13%
27	1	0.13%
28	1	0.13%
29	1	0.13%
30	2	0.27%
32	592	79.36%
33	8	1.07%
34	6	0.80%
35	35	4.69%
37	1	0.13%
40	7	0.94%
42	2	0.27%
43	1	0.13%
44	1	0.13%
45	1	0.13%
46	1	0.13%
47	1	0.13%
48	77	10.32%

The logo for IPv6, with the letters 'I', 'P', 'V', and '6' in a large, blue, 3D-style font. The letters are slightly offset and have a gradient effect.

# IPv6 Prefix Lengths



# IPv6 Network Capacity

- If IPv6 takes off, will there be capacity?
  - Has enough physical bandwidth been allocated to IPv6 networks, many of which are still in “test” phase.
  - Is there enough switching capacity through network devices. Some (older/smaller) boxes software switch IPv6. Some tunnel types are not well supported.

IPv6

# Supporting IPv6

- Do your IPv6 customers require support?
- Does your NOC speak IPv6?
- How many people understand IPv6?
- **What's your best single piece of advice on supporting an IPv6 network?**

IPv6

# IPv6 Future

- Increasing customer demand?
- IPv6 multicast anybody? No, seriously!
- Is Vista going to push IPv6?
  - Vista is IPv6 enabled out-of-the-box.
  - IPv6 tunnelling enabled automatically also.
  - Works even behind (most) NAT – Teredo.
  - Windows Meeting Space **requires** IPv6.

IPv6 » PNRPv2

# Vista & IPv6

```
C:\Windows\system32\cmd.exe
C:\Users\lab>ipconfig

Windows IP Configuration

Wireless LAN adapter Wireless Network Connection:

    Media State . . . . . : Media disconnected
    Connection-specific DNS Suffix . . . . . :

Ethernet adapter Local Area Connection:

    Connection-specific DNS Suffix . . . . . : eulab.net
    Link-local IPv6 Address . . . . . : fe80::87d1:2e66:5df3:acb4%8
    IPv4 Address. . . . . : 192.168.0.5
    Subnet Mask . . . . . : 255.255.255.0
    Default Gateway . . . . . : 192.168.0.1

Tunnel adapter Local Area Connection* 6:

    Connection-specific DNS Suffix . . . . . : eulab.net
    Link-local IPv6 Address . . . . . : fe80::5efe:192.168.0.5%16
    Default Gateway . . . . . :

Tunnel adapter Local Area Connection* 7:

    Media State . . . . . : Media disconnected
    Connection-specific DNS Suffix . . . . . :

Tunnel adapter Local Area Connection* 10:

    Connection-specific DNS Suffix . . . . . :
    IPv6 Address. . . . . : 2001:0:4136:e378:2c43:b56:2bfc:167d
    Link-local IPv6 Address . . . . . : fe80::87d1:2e66:5df3:acb4%10
    Default Gateway . . . . . :
```

- Really is IPv6 out of the box!
- 6to4 and Teredo tunnelling automatic!
- No user interaction or notification.
- Do you use NetMeeting?


# IPv6

# IPv6 Routing Table Size

- IPv6 routing table shows steady growth

- Drop in 2006 due to 6bone being ceased.
- Appears to be linear growth rather than exponential right now.
- Traffic growth still minimal.

# IPv6




Graph courtesy of Geoff Huston – <http://bgp.potaroo.net/>

NETWORK OPERATORS BOF – NANOG 39

# IPv6 Routing Table Size

- Number of unique ASes also growing in a linear fashion.

- Shows steady increase in IPv6 networks.
- Again, doesn't appear to be matched by traffic increase.



Graph courtesy of Geoff Huston – <http://bgp.potaroo.net/>


# IPv6



# IPv6 Routing Table Size

- Compares to the IPv4 routing table:
  - At what point does this stop increasing like this, if at all?
  - At what point does this become a problem?
  - Many devices have limited space for an IPv6 routing table. Wont be able to hold both IPv4 and IPv6 tables in the future.

# IPv6



Graph courtesy of Geoff Huston – <http://baq.potaroo.net/>

# IPv6 Routing Table Size

- When does this become a problem?
  - 800 prefixes doesn't take much space.
  - However, some vendors TCAM profiles don't allow a great deal for IPv6.
- Should we even worry about it?
- What should we do to keep this sane?

IPv6

# IPv6 Multi-homing

- Any further **sensible** thoughts on this?
- Seems that most people are just getting their own space right now.
  - How long is that going to work for?
- Without consensus on a better method, seems to be following the IPv4 path.

IPv6

# Questions, Comments & Other Issues

IPv6

NETWORK OPERATORS BOF – NANOG 39