

# IPv6 EYOD Experiments

Danny McPherson [danny@arbor.net](mailto:danny@arbor.net)

Randy Bush [randy@psg.com](mailto:randy@psg.com)



# Goal of Experiments

- Eat our own dog food
- Is IPv6 practically usable?
- What improvements need to be made to make it practically usable?
- NANOG, APRICOT & IETF all recently conducted experiments to sample and foster usability
- AFNOG, RIPE, ARIN, etc.. Experimenting as well
- Goals more educational and documentary than demonstrative

# The Enterprise Usage Model

- Consumers depends on ISPs, many models there
- Enterprises likely to be first ones who can NOT obtain new or additional IPv4 address space
- Enterprises are more “controlled” environments
- They’re Randy’s customers :-)

So what is it like to be an  
Enterprise that can only get  
IPv6 space?

# A Pure IPv6 Network Today...

- Would not get access to >99.9% of the 'net
- So you could see the dancing kame turtle (<http://www.kame.net>)

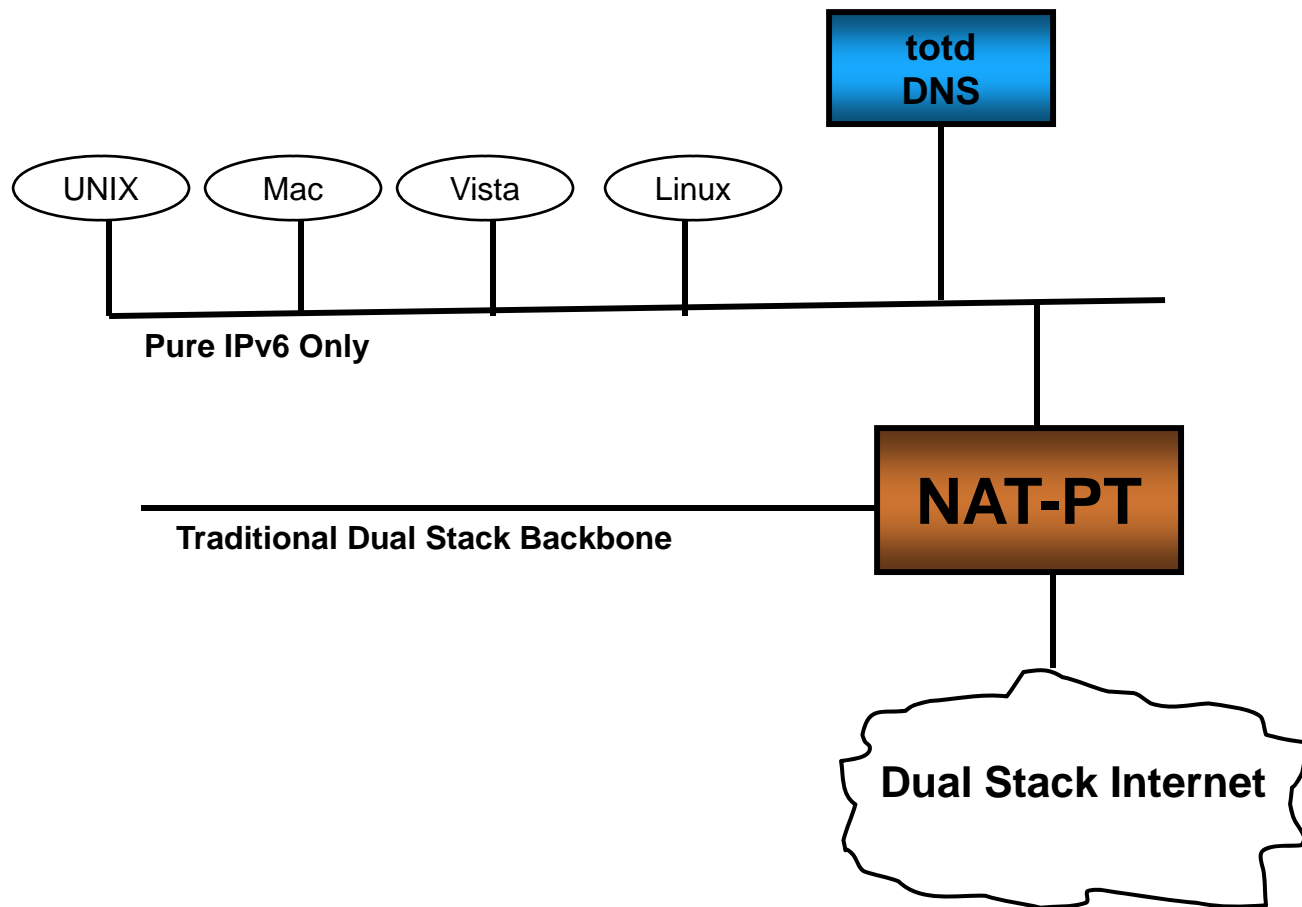
QuickTime™ and a  
decompressor  
are needed to see this picture.

- But not much else....
- Mainly because of bits on the wire incompatibilities with IPv4 and IPv6

# A Pragmatic IPv6 LAN Today

- An IPv6-only network with NAT-PT at the border and totd
- Now we can get to the IPv6 world AND the IPv4 Internet
- See the kame and see Yahoo!
- This should work for UNIX, Linux, Mac, Vista, ...

# NANOG/APRICOT IPv6 Network





# NAT-PT

- IPv6-only inside, dual-stack IPv4 and IPv6 on the other
- Stateful translation of TCP when no IP addresses are embedded
- ALGs needed for FTP, RTP, ... where addresses are embedded

# The DNS Hack

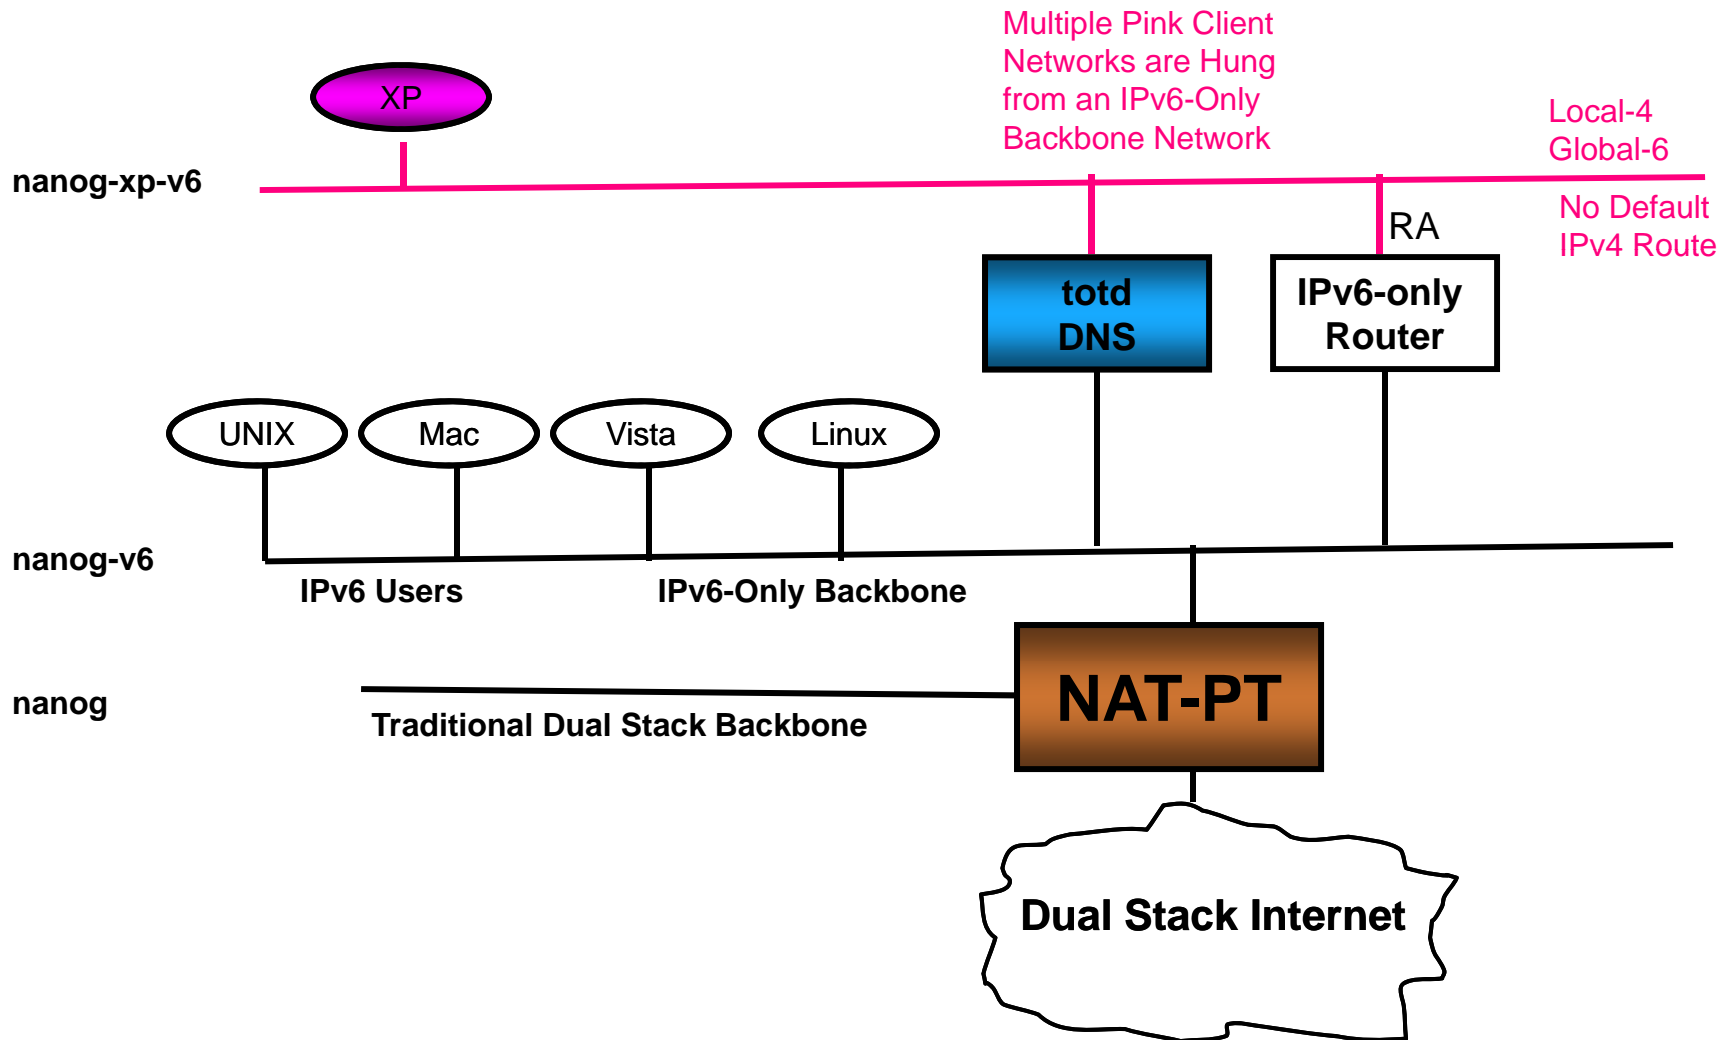
- On a pure IPv6 network, if I get an A record, what do I do?
- Panic, you can't use an A record!
- So the local DNS cache has a hack, *totd*, which takes an A, embeds it within a hacked IPv6 prefix, and synthesizes an AAAA
- NAT-PT knows the hack prefix, and strips it back to IPv4 to dual-stack



# Windows XP

- XP can move payload over IPv6
- But does NOT do DNS over IPv6 transport (nor AD, or other things! “Buy Vista”)
- So, the LAN has 1918 IPv4 space to carry DNS, but no exit for IPv4
- You get an IPv6 and IPv4 address, but should use the IPv6 for all real transport as IPv4 has no default

# Modified Network



# Access


- All SSIDs up from start
  - **nanog** normally fully dual-stack
  - **nanog-v6** with NAT-PT
  - **nanog-xp-v6** with 1918 DNS translation
- At APRICOT, for 90 minutes on Wednesday **apricot** went away, at NANOG for 60 minutes on Tuesday
- At NANOG some stats were gathered, none really collected at APRICOT

# What We Learned

- 3/4 of users said they could not get on net but it was actually more like half
- Every component except UNIXes had bugs: NAT-PT, Vista, MacOS, XP, ...
- The prize to MacOS, which dropped capital 'A' from DNS server entry
- We hope vendors now working to fix

# NANOG MAC Address Stats

MAC address counts by VLAN at NANOG42; Data from Tuesday, Feb. 19 2008  
(times corrected to PST)





# IETF 71 Experiment

- Held during IESG plenary session on Wednesday evening
- Already learned much from NANOG and APRICOT experiments
- Unique things:

# Other IETF observations..

- Some issues..
  - No workie:
    - iTunes Store
    - Apple Bug Reporter
    - Skype
  - Jabber (server issues with OpenFire per html v. txt transcripts, Adium client seemed OK - though many IM servers not reachable - e.g., google talk - though translators made available)
  - iChat will only work if there's still an IPv4 address on the system somewhere
  - [www.ripe.net](http://www.ripe.net) worked, but not others such as [lirportal.ripe.net](http://lirportal.ripe.net) and [ris.ripe.net](http://ris.ripe.net)
  - [www.apple.com](http://www.apple.com)
  - [jabber.psg.com](http://jabber.psg.com) had no AAAA records, although there is a [www.psg.com](http://www.psg.com) one. Exists now :-)

# IETF Experiment Notables

- There was a power outage just before the IPv6 outage
- `ipv6.google.com` launched
  - Local stuff worked
  - Not much else
- Mark Andrews provided BIND fix for Windows XP to subvert DNS resolution issue (i.e., run local bind resolver)
- All IETF Web sites worked
- Made /. - “The Night the IETF Shut Off IPv4”
- /. not reach via IPv6

# Double NAT - v4v6v4

- Double NAT model to accommodate
- Described in
  - draft-durand-v6ops-natv4v6v4-01.txt
- Distributed NAT model where globally routed v4 addresses are shared among a pool of users
- MTU & ALG considerations
- Still very much research

# How You Can Play

- <http://www.civil-tongue.net/6and4/>
- Has instructions for platforms
- Has instructions how to turn it back off when you go home!
- Is a wiki, so .. If you want to play:
  - Create a login ID, open a ticket and request editor privileges
  - You can also use tickets to report problems, send feedback, etc..
  - Share your experiences

# Thanks To....

- Chung-Hwa Telecom, MERIT & Comcast
- Philip Smith, Lucy Lynch, Joel Jaeggli
- IJ & ISOC
- Cisco & Philip Smith
- APRICOT, NANOG & IETF Sponsors
- And Philip Smith...