



IPv6 Transition & Operational Reality


IEPG / Chicago

2007.07.22


Randy Bush <randy@psg.com>

<http://rip.psg.com/~randy/070722.v6-op-reality.pdf>

What Should Have Happened



What Is Happening?



Why Is This Happening?

No transition plan

Declared victory before the hard part started

No real long term plan

No realistic estimation of costs

No support for the folk on the front lines

Victory will be next month

This Describes:

a - The invasion of Iraq

b - IPv6

c - DNSSec

d - All of the above

Myth: IPv4 is Running Out

- IPv4 Free Pool run-out in a few years
- This is almost exactly in line with the graphs of Frank Solensky ten years ago
- IPv4 will go to a *Trading Model*
- Registries will become *Title Agents*, not allocators, of IPv4 space
- ARIN developing full multi-RIR/LIR open source software to certify and verify title to IPv4 and IPv6 resources

Myth: IPv6 Transition is Easy

- IPv6 was designed with no serious thought to operational transition
- IPv6 is on-the-wire incompatible with IPv4
- Could have been avoided, e.g. if IPv6 had variable length addressing, IPv4 could have become the 32 bit variant
- There are no simple, useful, scalable translation or transition mechanisms

Myth: IPv6 Eliminates NATs

- An IPv6-only site can not reach IPv4 Internet because it can not source packets from an IPv4 address
- There will be significant IPv4-only Internet for a decade or more
- All IPv6 sites will need IPv4 space and will have NATs with ALGs
- IPv6 increases NAT use in short and medium term, i.e. a decade or more

Myth: IPv6 Reduces Routing Load

- Multi-homing in IPv6 is the same as in IPv4, there is no new routing model
- Traffic engineering in IPv6 is the same as in IPv4, no new TE model
- Enterprises will slice and dice their IPv6 /32s to handle branches etc.
- The routing table will fragment more and more over time

Myth: Transition Eases Routing

- One possible result is market-based BGP advertisements
- Operationally complex; routing is global so how are settlements distributed?
- This could push back on fragmentation
- But how much can an announcement cost to be less than the cost of running without any IPv4 compatibility

Myth: IPv6 Space is Infinite

- 64 bits goes to every LAN
- This leaves half the bits gone!
- Some folk use /64 for Point-to-Point!
- RIRs are giving away /32s
- In 15 years we will think of these as we now think of legacy /8s in IPv4 space

Myth: IPv6 has Better Security

- IPv6 does nothing IPv4 does not, though it promised to
- IPSec is the recipe in either case
- IPSec does not work well in a mixed IPv4/IPv6 environment (think VPN from an IPv4-only hotel room)
- It is true that address space scanning will be somewhat harder
- Ha Ha, think botnet scanning and a black market in hot space

Myth: Incremental Deployment

- For an enterprise, the entire chain, from database back end, through applications, through firewalls, to the border router must all support v6 or the enterprise can not deploy
- For ISP, provisioning systems, monitoring, measurement, billing, ...
- And everyone needs support from all their vendors

Myth: Routers Support IPv6

- But not 100% in hardware
- Especially not if you add ACLs
- And all vendors are not spinning the ASICs to solve this
- Not all v4 features are supported over IPv6: MIBs, SNMP over v6, ...

Myth: No Static Numbering

- IPv6 Auto Configuration is not widely used in enterprise as security policy prefers known (i.e. DHCP) addresses
- Similarly, ISP backbone addresses and customer addresses must be known for logging, audit, CALEA, ...

Myth: IPv6 is Deployed

- Pioneers are still moving cautiously
- Early adopters are just starting to enter the game
- Actual measured traffic is very small (so it makes routers look as if they can handle the traffic)
- But there are good anecdotes


Myth: IPv6 Will Replace IPv4

- Not given current lack of **universal** vendor support from back end to border router
- It is far easier to use NAT and IPv4
- IPv4 with NATs requires no new expense, conversion, training, ...
- This is architecturally horrible, it is just financial reality

The Reality

- "96 more bits, no magic"
-- Gaurab Raj Upadhaya
- But we sure need more bits!
- The key questions are how to use them?
- How to transition without losing anyone or anything?

What Can We Do?



How?

- Identify current transition problems
- See that they are fixed
- Ask the IETF to fix the outstanding protocol issues
- Push vendors to support IPv6 and the tools for us to transition
- Registries need to prepare to issue titles to IPv4 and IPv6 space

What We Should Not Do

- Pretend that there are no transition problems. It just makes things harder.
- Give away IPv6 space in strange ways to "promote" IPv6. IPv4 run-out will promote IPv6 for us.
- Make messes we will have to live with forever.

Areas of Concern/Study

- Global Issues
- Administrative Infrastructure
- Layers 1 and 2
- Backbone Engineering
- Last Mile/Kilometer
- Consumer/SOHO Self-Installed CPE
- Enterprise
- Server Farm
- Campus
- Exchange Points
- Applications
- Telephony
- More?

Global Issues

- How does a user at a v6-only site get to the Internet, i.e. a v4-only site?

They don't!

- What can be done to help as much as possible?

Administrative Infrastructure

- DNS
 - BIND 9 seems to fully support IPv6
 - Registrars need to support delegation to IPv6 nameservers
 - Registrars need to support IPv6 glue records
- RIRs - ARIN developing open source package for X.509 certification of resource 'ownership'

Layers 1 and 2

- DOCSIS 3.0 for Cable
 - MTU limit of 1518
 - CMTS support lacking
- 802
 - All media protocols support IPv6
 - While the protocols support IPv6, this does not at all mean that implementations do

Backbone Engineering

- Core Routing - conversion to dual stack is slow
- IS-IS support good, OSPFv3 weaker
- Provisioning, Address Assignment, DNS, ...
- DHCPv6 and DNS Integration
- Monitoring and Measurement over v6?
- New line cards are often required!

The Last Kilometer

- Authentication and session setup, e.g. PPPoE, IPoE, DHCP
- Provisioning, back-end database, ...
- "How to scale the routing/provisioning combo to deal with million of customers using stable prefix delegation?"

Consumer Self-Installed CPE

- \$50 DSL Modems do not support v6
- \$50 Firewalls do not support v6
- Teredo does not really scale
- shim6 is not deployable due to security and routing model issues

Enterprise

- Databases, PeopleSoft, Siebold, Business Applications, ...
- Firewalls, VPNs, Access, ...
- Millions of lines of in-house code
- NFS Appliances, unknown
- If one link in the business application production chain is not there, it does not transition

Server Farm

- Lack of wide availability of V6 transit, especially as a non-beta service
- Very limited selection of hardware load balancers which support IPv6
- Storage Networks? (unknown)
- Back-end systems need reworking
- Akamai and other CDN providers

Campus

- Apache is OK, Squid is not
- Wireless provisioning is weak
- Perl `inet_ntoa(inet_aton($ip))` does not work
- EUI-64 makes some host based access control very complicated
- P2P Applications, status unknown

Exchange Point

- Management & Measurement
- NTP and other IX services

Applications

- Where is the web page with a matrix of application by platform showing which are v6 capable and clickable link on how to turn it on?
- Many applications which support v6 have sufficiently poor performance that early adopters are being told to turn v6 off
- XP will not work in a v6-only environment, because it does not support DNS queries over IPv6

Telephony

- SIP inter-networking between IPv6 only UE and IPv4 only UE
- SIP, RTP, H323, NAT traversal,
...

Why is Japan in Better Shape?

- Folk with vision (i.e. Murai) convinced the government that early movement to IPv6 was wise for Japan
- Government \$support\$ IPv6 research
- Government \$support\$ IPv6 development by industry, vendors, ...
- Government give\$ tax incentive\$ to enterprises which become v6 compatible

How You Can Help

<http://www.civil-tongue.net/clusterf/>

write to randy@psg.com
if you can contribute

Please!

Thanks To

**Internet Society for Support of
Operational Transition Analysis**

ARIN

for Research Support

Internet Initiative Japan