

Internet Address Exhaustion

Geoff Huston
APNIC

The Roots of Open Systems

- Unix
- TCP/IP

Both technologies benefited from open source reference implementations that allowed derivative development while maintaining a common core of functionality

TCP/IP has effectively displaced most proprietary competitors over time

Openness is Good!

Open Technologies are:

- Generally Accessible
- Competitively Neutral
- Functionally Extensible
- Commercially Exploitable

Openness has been one of the main strengths behind the ascendancy of the Internet

So What?

Being “open” is one thing

Staying “open” is another

What is important here is that the key task for the Internet is **sustaining** openness

Why does Openness require constant attention?

Useful technologies are rarely static:

- Technology evolves
- Uses change
- Exploitation models change

There is a need to continuously define the line of demarcation between:

- competitive interests and private "ownership" and
- common interest and "public good"

Challenges to Open Networking

- Net Neutrality
- Next Generation Networks
- Mobility and Mobile service evolution
- Triple/Quad-play schemes

Let's focus our attention just a little...

What is the most critical issue in today's Internet that is going to shape the entire future of open networking?




We've run out of addresses!




oops!

Inexorable Growth


IPv4 Global Address Allocations



Inexorable Growth




Inexorable Accelerating Growth




Current Status of IPv4

IPv4 Address Pool Status



We had a plan ...




The Theory


The idea was that we would never “run out” of IPv4 addresses

- Industry would see the impending depletion and gradually and seamlessly fold IPv6 into their products and services
- We would be an all-IPv6 Internet before we ever had to use the last IPv4 address
- And no customer would see any change during the entire process

Testing the Theory: Tracking IPv4



IPv4 Exhaustion




IANA Exhaustion February 2011

First RIR Exhaustion
July 2011



Where are we with IPv6 deployment?



©2011 Google


<http://www.google.com/intl/en/ipv6/statistics/>

Something's not right!


- We were meant to have completed the transition to IPv6 BEFORE we completely exhausted the supply channels of IPv4 addresses



The IPv6 Transition Plan - V2




Is this Plan Feasible?




Deploy IPv6 across some 1.8 billion users, with more than a billion end hosts.

Is this Plan Feasible?




Deploy IPv6 across some 1.8 billion users, with more than a billion end hosts, and upgrade hundreds of millions of routers, firewalls and middleware units.

Is this Plan Feasible?




Deploy IPv6 across some 1.8 billion users, with more than a billion end hosts, and upgrade hundreds of millions of routers, firewalls and middleware units, and audit billions of lines of configuration codes and filters.

Is this Plan Feasible?



Deploy IPv6 across some 1.8 billion users, with more than a billion end hosts, and upgrade hundreds of millions of routers, firewalls and middleware units, and audit billions of lines of configuration codes and filters, and audit hundreds of millions of ancillary support systems.


Is this Plan Feasible?




Deploy IPv6 across some 1.8 billion users, with more than a billion end hosts, and upgrade hundreds of millions of routers, firewalls and middleware units, and audit billions of lines of configuration codes and filters, and audit hundreds of millions of ancillary support systems -

all within the next 200 days.

The IPv6 Transition Plan - V3



What then?



- Some possible scenarios to sustain a growth rate of $> 250\text{M}$ new services every year:
 - Persist in IPv4 networks using more NATs
 - Transition to IPv6

IPv4 NATs Today



- Today NATS are largely an externalized cost for ISPs
 - Customers buy and operate NATS
 - Applications are tuned to single-level NAT traversal
 - Static public addresses typically attract a tariff premium in the retail market
 - For retail customers, IP addresses already have a market price!

The “Just Add More NATs” Option



- Demand for increasing NAT “intensity”
 - Shift ISP infrastructure to private address realms
 - Multi-level NAT deployments both at the customer edge and within the ISP network
 - This poses issues in terms of application discovery and adaptation to NAT behaviours
- Market cost for public addresses will increase to reflect realities of scarcity and higher exploitative value

NAT Futures



- NATs represent just more of the same
 - NATs are already extensively deployed today
- But maybe not...
 - More intense use of NATs will alter the network's current architectural model, as ports become the next scarce shared resource
 - Applications must change to reflect an ever smaller aperture through which the Internet can be seen and used
 - Increasing cost will be pushed back to consumers as price escalation

NAT Futures



- How far can NATs scale?
 - Not well known, but the unit cost increases with volume
 - What are the critical resources here?
 - NAT binding capacity and state maintenance
 - NAT packet throughput
 - Private address pool sizes
 - Application complexity
 - Public Address availability and cost

NAT Futures




- ?
- In the escalating complexity curve, when does IPv6 get to look like a long term cheaper outcome?

The Other Option: IPv6





The Other Option: IPv6



- Transition to IPv6
 - But IPv6 is not backward compatible with IPv4 on the wire
 - So the plan is that we need to run some form of a “dual stack” transition process
 - Either dual stack in the host, or dual stack via protocol translating proxies


Dual Stack Transition to IPv6



Theology— Phase 1

- “Initial” Dual Stack deployment:
 - Dual stack networks with V6 / V4 connectivity*
 - Dual Stack hosts attempt V6 connection, and use V4 as a fallback


Dual Stack Transition to IPv6



Theology – Phase 2

- “Intermediate”
 - Older V4 only networks are retro-fitted with dual stack V6 support


Dual Stack Transition to IPv6



Theology - The final outcome

- “Completion”
 - V4 shutdown occurs in a number of networks
 - Connectivity with the residual V4 islands via DNS ALG + NAT-Protocol Translation
 - Outside the residual legacy deployments the network is single protocol V6

Dual Stack Assumptions




- That we could drive the entire transition to IPv6 while there were still ample IPv4 addresses to sustain the entire network and its growth
- Transition would take some (optimistically) small number of years to complete
- Transition would be driven by individual local decisions to deploy dual stack support
- The *entire* transition would complete *before* the IPv4 unallocated pool was exhausted

Dual Stack



- Dual Stack transition is not a “or” proposition
 - Its not a case of IPv4 today, IPv6 tomorrow
- Dual Stack transition is an “and” proposition
 - It’ s a case of IPv4 **AND** IPv6
- But we don’ t know for how long
 - So we need to stretch IPv4 out to encompass tomorrow’ s Internet, and the day after, and ...

Implications




- Whether its just IPv4 NATs OR transition to IPv6 ...
 - IPv4 addresses will continue to be in demand far beyond the date of exhaustion of the unallocated pool
 - In the transition environment, all new and expanding network deployments will need IPv4 service access and addresses for as long as we are in this dual track transition
- But the process is no longer directly controlled through today' s address allocation policies
 - the IPv4 address pool in the sky will run out!
 - the mechanisms of management of the IPv4 address distribution and registration function will necessarily change

Making IPv4 Last Longer



- Its not the IPv4 address pool that's fully consumed
 - It's the unallocated address pool that's been consumed
 - 20% of the address space is not advertised in global routing
 - widespread use of NATs would yield improved address utilization efficiencies
- So we could “buy” some deviant Second Life for IPv4
 - But it won't be life as we've known it!
 - It will be predicated on the operation of a market in IPv4 addresses
 - And such a market in addresses will not necessarily be open, accessible, efficient, regulated or even uniformly visible
 - This prospect is more than a little worrisome

Making IPv4 Last Longer



- Some ideas I've observed so far:
 - Encourage NAT deployment
 - Larger Private Use Address Pool
 - Policies of rationing the remaining IPv4 space
 - Undertake efforts of IPv4 Reclamation
 - Deregulate Address Transfers
 - Regulate Address Transfers
 - Facilitate Address Markets
 - Resist Address Markets

Making IPv4 Last Longer



- For how long?
- For what cumulative address demand?
- For what level of fairness of access?
- At what cost?
- For whom?
- To what end?
- What if we actually achieve something different?
 - How would the Law of Unintended Consequences apply here?
 - Would this negate the entire “IPv6 is the solution” philosophy?

Why is the Internet wedged on IPv6?

Why is the Internet wedged on IPv6?

Cost and Benefit are not aligned:

Those folk who have been historical "losers" with the open Internet - the carriers - are now being asked to make the bulk of the investment in IPv6 access infrastructure

They are understandably reluctant to make further investments that in the end just worsen their long term revenue prospects

Why is the Internet wedged on IPv6?

Cost and Benefit are not aligned:

Those folk who have benefitted from an open network in the past are increasingly ambivalent about open networks and IPv6 into the future

They are now part of the set of entrenched incumbents, with market positions to protect

Why is the Internet wedged on IPv6?

Cost and Benefit are not aligned:

Those folk with the most to gain in the longer term from continuation of an open network - consumers - do not necessarily act from day to day in their own long term interests

i.e. consumers are unwilling to fund this transition through higher prices for Internet services

Is the Open Pendulum swinging back?

Continued delay by incumbents to embrace IPv6 allows further consolidation, and increased ability by incumbents to define (and limit) the parameters of future competition.

What is at risk here is a truly open network infrastructure

If we really want to keep a truly Open Internet...

Then we need to alter our environment to favor the rapid adoption of IPv6!

Figuring out how to do this is proving to be extremely challenging

Thank You



STOP
ALL DRIVERS
WAIT HERE
FOR INSTRUCTIONS

