The IPv6 Condition

Geoff Huston
Chief Scientist
APNIC

IPv6 Evolution?

- The Internet as an evolving lifeform or ecosystem
  - If IPv6 can offer clearly superior value propositions to the industry then it will be deployed
    - The “invisible hand” of competitive market forces will lead the industry to adopt IPv6 naturally
  - Inferior technologies will wither away as they cease to offer any utility or lasting value
- Just let nature (the market) take its course!
  - though result is undesigned and unpredictable, should not be viewed as decay. Its adaptation.

Is IPv6 really evolutionary?

Or, to use a multi-choice variant of this question: Is an industry-wide IPv6 transition going to proceed as:

- **extinction** - acting as a catalyst to take a step to some other entirely different technology platform that may have little in common with the Internet architecture as we understood it?
- **evolution** - by migrating existing IPv4 networks and their associated service market into IPv6 in a piecemeal fashion?
- **revolution** - by opening up new service markets with IPv6 that directly compete with IPv4 for overall market share?

Extinction?

- The original IP architecture is dying – if not already terminally dead
  - Coherent transparent end-to-end is disappearing
  - Any popular application today has to be able to negotiate through NATs, ALGs and other middleware
  - Peer-to-peer networks now require mediators and agents (SpeakFreely vs Skype), plus stun, ice,…
  - Efforts to impose overlay topologies, tunnels, virtual circuits, traffic engineering, fast reroutes, protection switches, selective QoS, policy-based switching on IP networks appear to have simply added to the cost and detracted from the end user utility
- It was a neat idea, but we killed it!

IPv6?

- We’ve all heard views that:
  - It represents a very marginal change in terms of design decisions from IPv4
  - It did not manage to tackle the larger issues of overloaded address semantics
  - It did nothing to address routing scaling issues
  - And the address architecture is so broken that it yields just 48 useful bits out of 128

Alternate Worlds?

- Is there anything else around?

  Nope - not in the near term

- How long would a new design effort take?

  Tough – At least a decade or longer

  (we're not getting any smarter!)

- Would an entirely new design effort end up as a marginal outcome effort – would we be looking at no more than a slightly different set of design trade-offs within a common set of constraints?

  Probably

  (all that effort to get nowhere different!)

IPv6?

- **extinction** - acting as a catalyst to take a step to some other entirely different technology platform that may have little in common with the Internet architecture as we understood it?
- **evolution** - by migrating existing IPv4 networks and their associated service market into IPv6 in a piecemeal fashion?
- **revolution** - by opening up new service markets with IPv6 that directly compete with IPv4 for overall market share?

So should we evolve?

- The general answer appears to be "**yes**" for most values of "**we**"
- The possible motivations differ for each player:
  - Allow for networks with more directly addressed end points
  - Reduce per-address cost
  - Reduce application complexity
  - Increase application diversity and capability
  - Allow direct peer-to-peer networking
  - Allow utility device deployment
  - Leverage further efficiencies in communications

Pressure for Change?

- The pain of deployment complexity is not shared uniformly:
  - ISPs are not application authors -- thank god!
  - ISPs are not device manufacturers -- also a good thing!
- There appear to be no clear “early adopter” rewards for IPv6
  - Existing players have strong motivations to defer expenditure decisions
  - New players have no compelling motivations to leap too far ahead of their seed capital and customer base
  - All players see no short term incremental benefit in early adoption
  - And many players short term interests lie in deferral of additional expenditure
  - The return on investment in the IPv6 business case is simply not evident in today’s ISP industry

When?

- So the industry response to IPv6 deployment appears to be:

  “yes, of course, but **later**”

What is the trigger for change?

- At what point, and under what conditions, does a common position of “**later**” become a common position of “**now**”?

- So far we have no clear answer from industry on this question

The Case for IPv6

- IPv4 address scarcity is already driving network service provision.
  - Network designs are based on address scarcity
  - Application designs are based on address scarcity
- We can probably support cheaper networks and more capable applications in networks that support clear and coherent end-to-end packet transit
- IPv6 is a conservative, well-tested technology
- IPv6 has already achieved network deployment, end host deployment, and fielded application support

- For the Internet industry this should be a **when** not **if** question

But….

- But we are not sending the right signals that this is ‘cooked and ready’ - we are still playing with:
  - The Address Plan
  - Aspects of Stateless auto-configuration
  - Flow Label
  - QoS
  - Security
  - Mobility
  - Multi-addressing
  - Multi-homing
  - Routing capabilities
  - Revisiting endpoint identity and network locator semantics

But…

- IPv6 is not cleanly “backward compatible”
  - Piecemeal deployment leads to piecemeal connectivity, leads to piecemeal services

But ….

- Dual stack in the host implies dual protocols in the network, and dual service presentations at the server end
  - This is neither cheap nor easy to deploy in a high volume robust manner

The Business Obstacles for IPv6

- Deployment by regulation or fiat has not worked in the past – repeatedly
  - GOSIP anyone?
- There are no network effects that drive differentials at the edge
  - its still email and still the web
- There is today a robust supply industry based on network complexity and address scarcity
  - And they are not going to go away quietly or quickly
- The prospect of further revenue erosion for network service providers

More Business Obstacles for IPv6

- Having already reinvested large sums in packet-based data communications over the past decade there is little investor interest in still further infrastructure investment at present
  - The only money around these days is to fund MPLS fantasies!
- There is no current incremental revenue model to match incremental costs
  - Oops! Customer won't pay more for IPv6
- IPv6 promotion may have been too much too early – these days IPv6 may be seen as tired not wired
  - Too much powerpoint animation!
- Short term individual interests do not match long term common imperatives
  - The market response is never an intelligent one
- "Everything over HTTP" has proved far more viable than it should have

Meet the Enemies!

- “As easy as plugging in a NAT”
  - NATs are an excellent example of incremental deployment and incremental cost apportionment
- The search for perfection
  - Constant adjustment of the protocol specifications fuels a common level of perception that this is still immature technology
- The search for complexity
  - Pressure to include specific mechanisms for specific scenarios and functionality as a business survival model

The current situation

The entire Internet service portfolio appears to be collapsing into a small set of applications that are based on an even more limited set of HTTP transactions between servers and clients

This is independent of IPv4 or V6

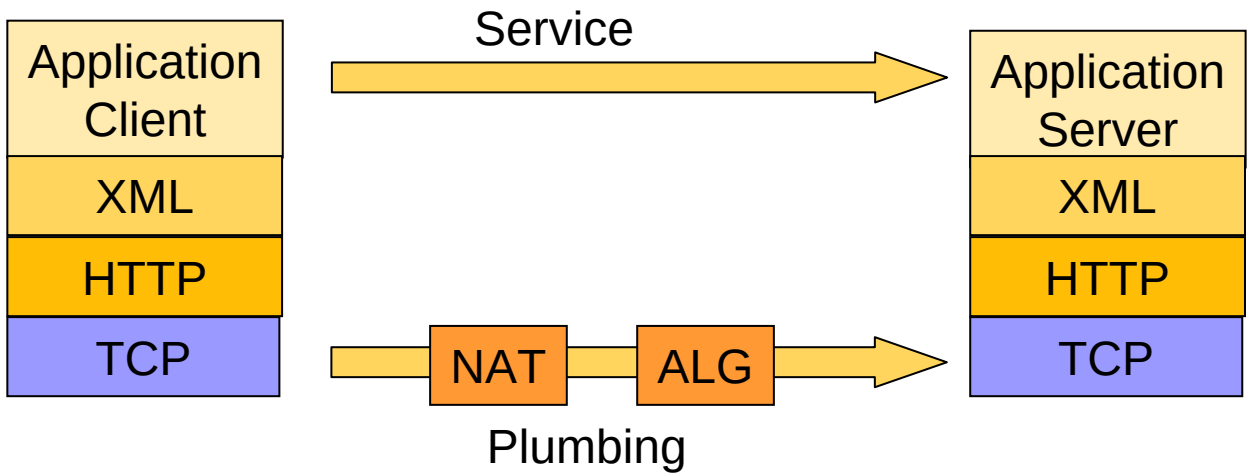

Maybe it’s just deregulation

- Near term business pressures simply support the case for further deferral of IPv6 infrastructure investment
- There is insufficient linkage between the added cost, complexity and fragility of NAT-based applications at the edge and the costs of infrastructure deployment of IPv6 in the middle
  - Deregulated markets are not perfect information markets – pain becomes isolated from potential remedy
  - Markets often cannot readily trade off short term cost against longer term benefit

IPv6?

- **extinction** - acting as a catalyst to take a step to some other entirely different technology platform that may have little in common with the Internet architecture as we understood it?
- **evolution** - by migrating existing IPv4 networks and their associated service market into IPv6 in a piecemeal fashion?
- **revolution** - by opening up new service markets with IPv6 that directly compete with IPv4 for overall market share?

Learning from IPv4

- IPv4 leveraged:
  - cheaper switching technologies
  - more efficient network use
  - lower operational costs
  - structural cost transferral
- IPv4 represented a compelling and revolutionary business case of stunningly cheaper and better services to end consumers, based on the silicon revolution

The IPv6 Condition

- There are no compelling technical feature levers in IPv6 that are will drive new investments in existing IP service platforms
- There are no compelling revenue levers in IPv6 that will drive new investments in existing IP service platforms

So why IPv6?

- IPv6 represents an opportunity to embrace the communications requirements of a device-dense world
  - more than PCs
  - Device population that is at least some 2 – 3 orders of magnitude larger than today's Internet

- BUT - Only if we can further reduce IP service costs by a further 2 -3 orders of magnitude
  - Think about prices of the level of $1 per DSL service equivalent per year

IPv6 - From PC to iPOD to iPOT

If we are seriously looking towards a world of billions of chattering devices then we need to look at an evolved communications service industry that understands the full implications of the words “**commodity**” and “**utility**”

The IPv6 Revolutionary Manifesto

- **Volume over Value**
  - Supporting a network infrastructure that can push down unit cost of packet delivery by orders of magnitude
  - Commodity volume economics can push the industry into providing
    - even “thicker” transmission systems
    - simpler, faster switching systems
    - utility-based provider industry
    - Lightweight application transaction models

- So it looks like the IPv6 future could be **revolution** where IPv6 is forced into direct customer competition with existing IPv4+NAT networks

- And the primary leverage here is one of **cheaper** and **bigger**, and not necessarily *better*

Maybe: …

We need to regard IPv6 in different terms:

Perhaps we should look at IPv6 as the enabler for **vastly** larger networks

And stop looking for higher value propositions with IPv6 networks, but instead look for lower costs in switching IPv6 packets

But…is this realistic?

- Is it really possible that there are further cost economies to be realized in the carrier IP network industry?

Where is the next bloat to strip?

- Transmission infrastructure?
  - Fibre optics vs Physics
  - Spread spectrum wireless vs spectrum pollution
  - ?
- Switching?
  - Electrical vs power and speed
  - Optical vs physics
- O&M?
- Cost of Finance?
- Investor returns?

- A vastly ‘cheaper’ network is unlikely in the near to medium term
  - Irrespective of volume drivers
- Which doesn’t look good for IPv6
- And it makes the “revolutionary” IPv6 approach of achieving vastly lower cost points though higher volumes for IPv6 look rather unsatisfactory as a viable outcome!

IPv6?

- **extinction** - acting as a catalyst to take a step to some other entirely different technology platform that may have little in common with the Internet architecture as we understood it?
- **evolution** - by migrating existing IPv4 networks and their associated service market into IPv6 in a piecemeal fashion?
- **revolution** - by opening up new service markets with IPv6 that directly compete with IPv4 for overall market share?

Then what's left?

- Making IPv4 + NATS work for ever?
  - Unlikely?
- Forced IPv6 conversion?
  - Unlikely?
- Something else?
  - Well nothing new – no time!

- So maybe all that's left is to tinker around in the protocol stack to see if there is any leverage to be gained here

Splitting Identity and Location?

- Hosts care deeply about absolute identity
- Networks care deeply only about relative location
  - All a network really cares about is to associate incoming packets with the relative location of the network exit point
  - After that its SEP!
- Its actually the identity component of IPv4 addressing that's under stress, not the network address component
  - And HIP and SHIM6 are both decent experimental prototypes of how these differing semantic address components can be split at the endpoint rather than within the network infrastructure elements

So, possibly:

- It need not be an IPv4 / IPv6 issue at all!
  - It could actually be about what element of end-to-end address semantics is essential at the host-to-host transport level and what part is devolved to a mapping / translation problem at the network level

Maybe the issue we face with IP tomorrow is really all about the fundamentals of networking architectures rather than the size of the address field in the packet header

Thank You