Internet Evolution and IPv6

IPv6 - the BGP view

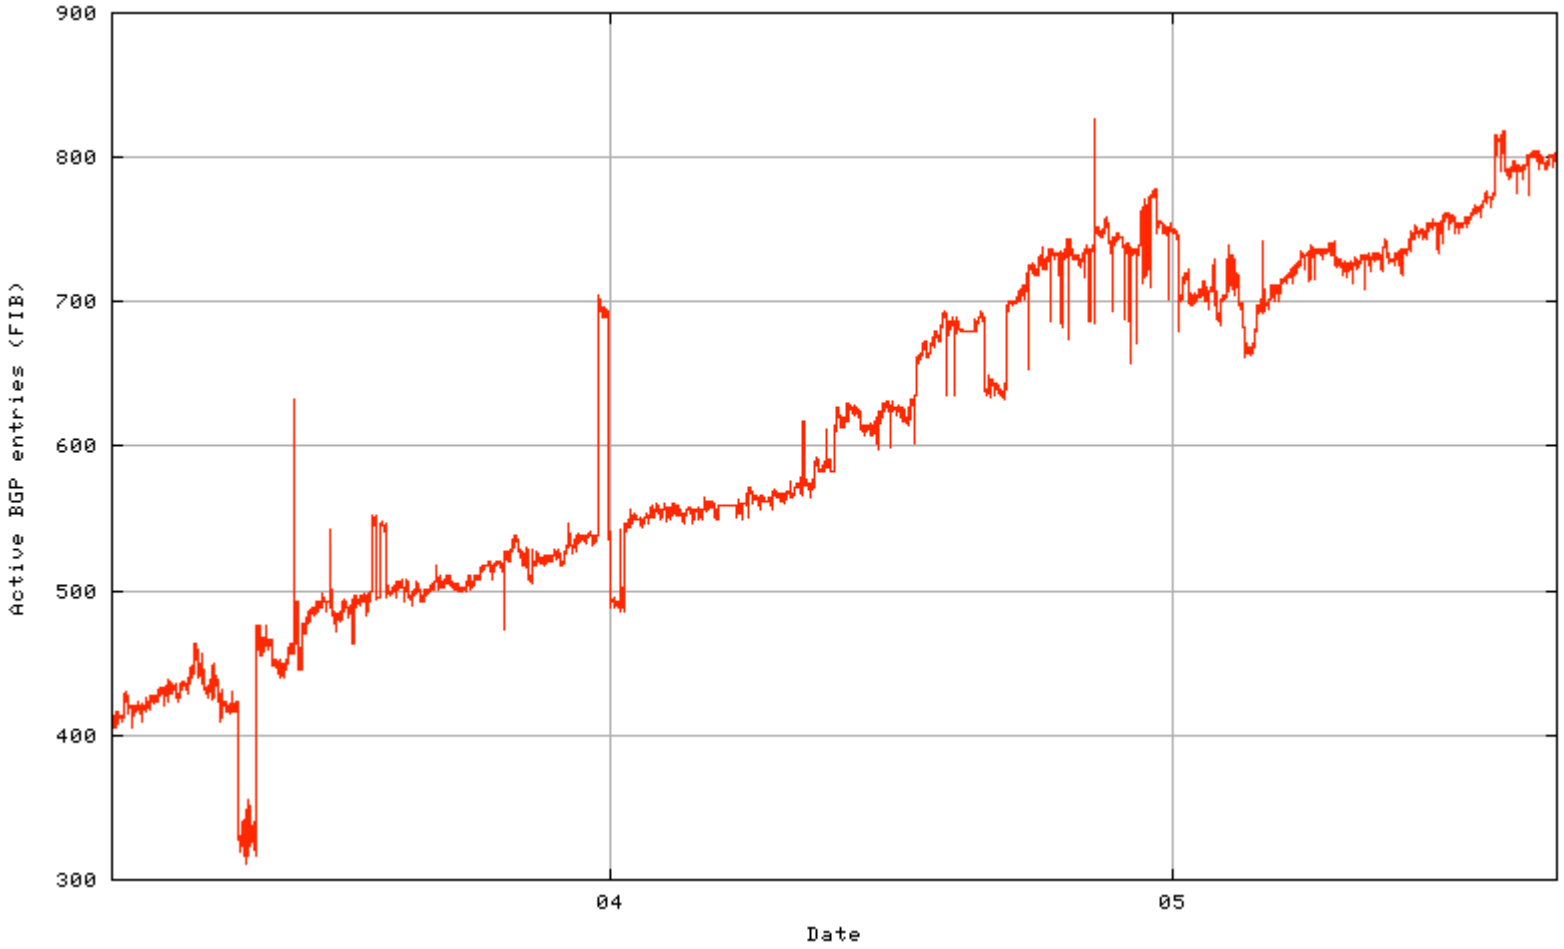

IPv4 – the BGP view

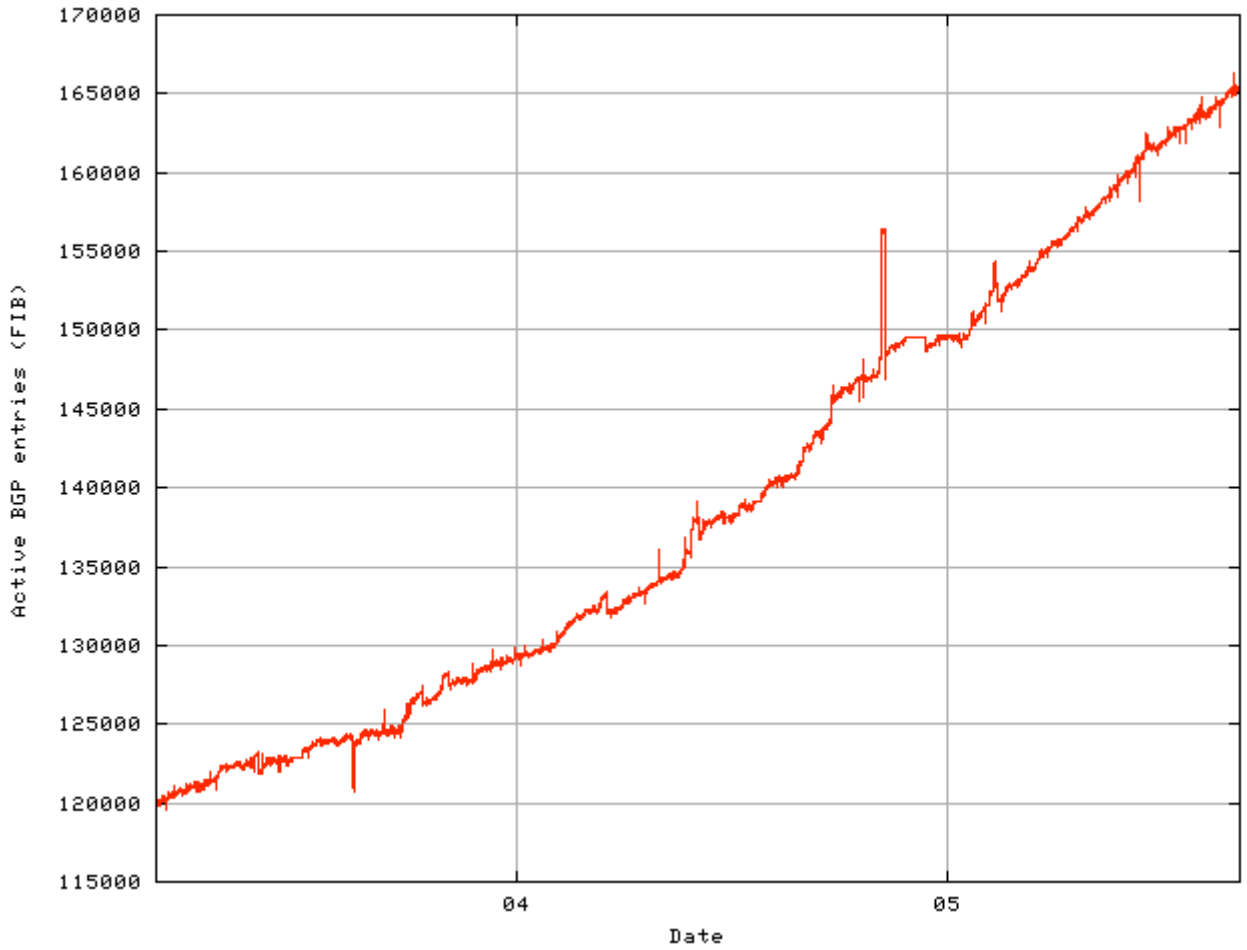

IPv6 Adoption – AS Count

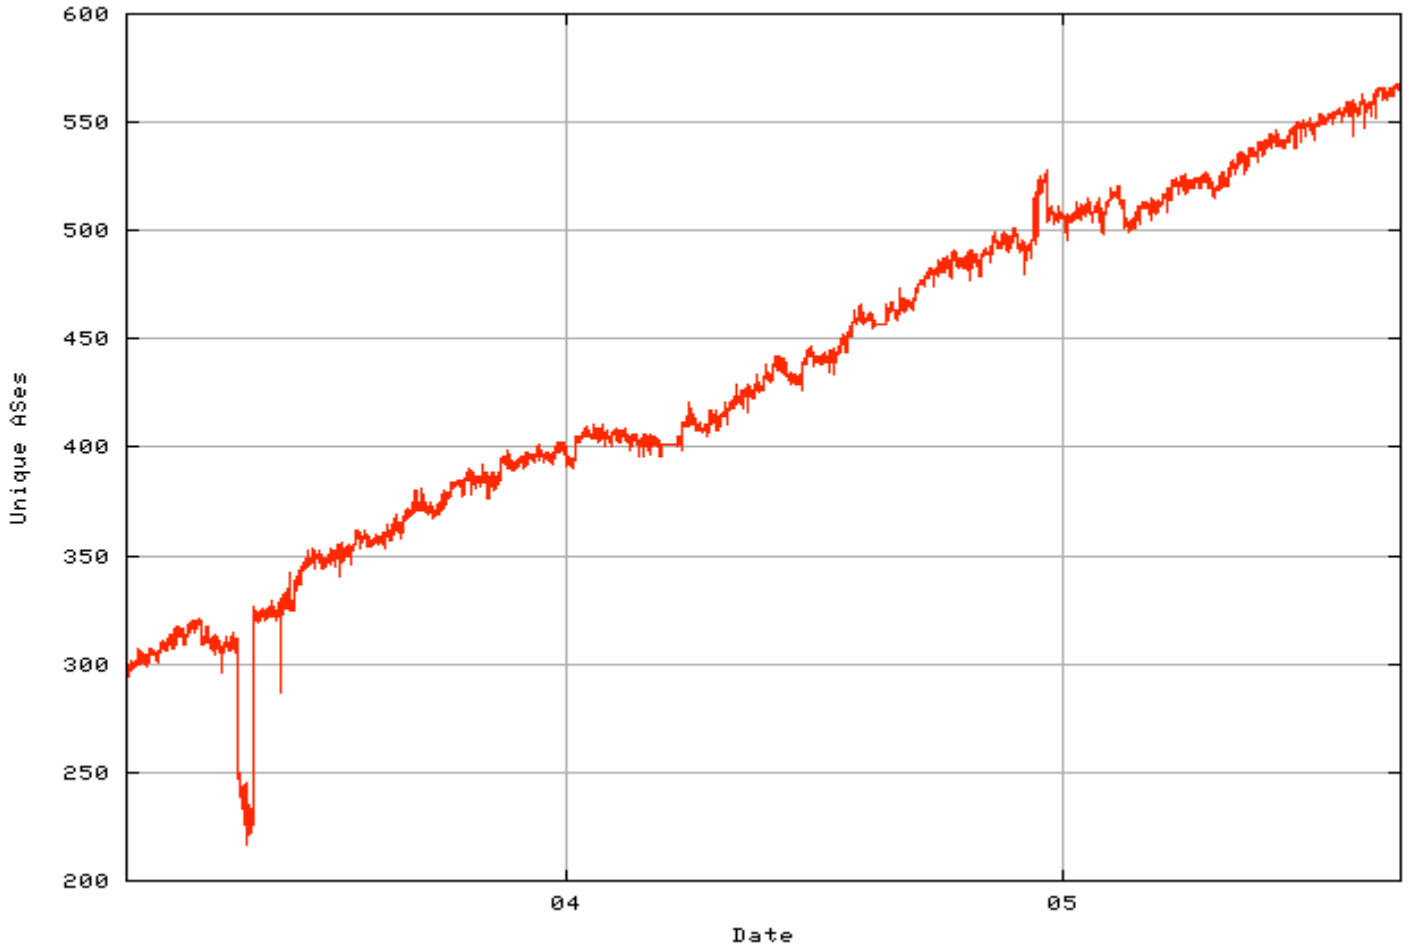

IPv4 Expansion – AS Count

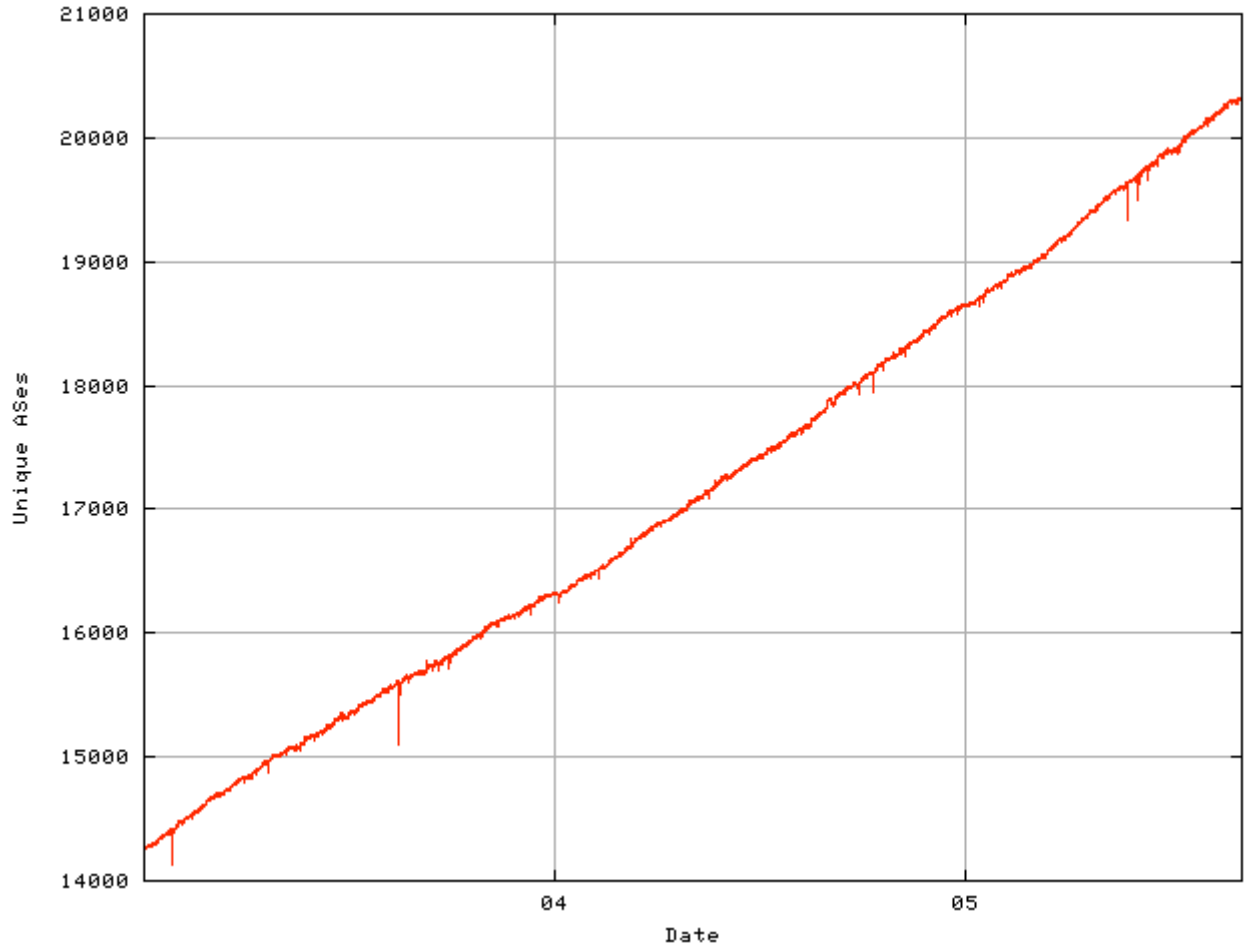

IPv6 vs IPv4 Rates – AS Count

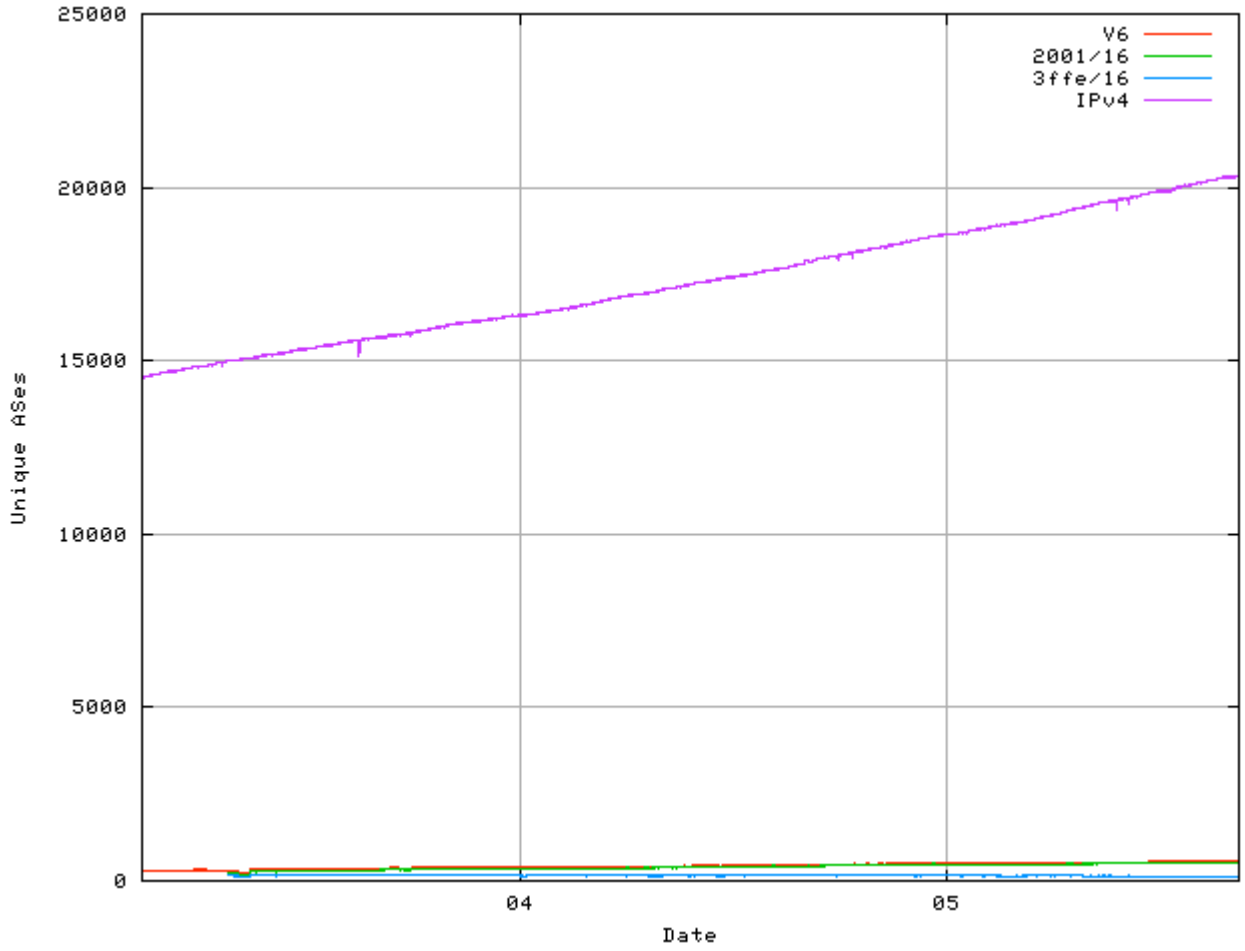

From Optimism to Conservatism

- We've learned that optimism alone is no substitute for knowledge and capability within this industry

- Current conservative period of consolidation rather than explosive growth
  - Investment programs need to show assured and competitively attractive financial returns across the life cycle of the program
  - Reduced investment risk implies reduced levels of innovation and experimentation in service models
  - Attempts to combine communications with additional services to create value-added service bundles
  - Accompanied by greater emphasis of financial returns from existing infrastructure investments

IPv6 - some industry options

Is an industry-wide IPv6 transition going to proceed as:

- **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 compete with IPv4 for overall market share?
- **extinction** act as a catalyst to take a step to some other entirely different technology platform for communications that has little in common with the Internet architecture as we understood it?

What is the story with IPv4?

- The original IP architecture is dying
  - 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)
  - 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

IPv4 address depletion?

One View: We effectively ran out of IPv4 addresses at the edge of the network at the time when NAT deployment became prevalent

In today's retail environment one stable public IPv4 address can cost as much as megabit DSL access

We are running out of unallocated addresses to inject into the network

- that does not mean addresses will no longer be available
- it probably just means that the nature of the distribution function and the pricing function will change

Today

- We are engineering applications and services in an environment where NATs, Firewalls and ALGs are assumed to be part of the plumbing
  - Client-initiated transactions
  - Application-layer identities
  - Agents to orchestrate multi-party rendezvous
  - Multi-party shared NAT state
- All this complexity just results in more fragile applications

So should we move on?

- 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 for IPv4 NAT deployment is not shared uniformly:
  - ISPs are not application authors
  - Existing players have strong motivations to defer expenditure decisions
  - New players have no compelling motivations
  - Many players see no incremental benefit in early adoption
  - Many players short term interests lie in deferral of additional expenditure
- There appear to be no clear early adopter rewards for IPv6
  - The return on investment in the business case is simply not evident

When?

- So the industry response appears to be “**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

IPv6 or something else?

- Is there anything else around today that takes a different view how to multiplex a common communications bearer?
- How long would a new design effort take?
- Would an new design effort end up looking at an entirely different architecture? Or would it be taking a slightly different set of design trade-offs within a common set of constraints?

Packet Switching attributes

- Packet switching represents a weak form of control design, is harder to operate than circuits, and tends to push cost, value (and revenue) off the network and into the edge
- Packet switching is cheaper, is more efficient, is cheaper, is less constraining on service models, is cheaper, enables more edge innovation, and is cheaper

Common Constraints
Service Control Capabilities

- No communications network can intrinsically change human behaviour, nor can it provide robust ‘cures’ for spam, IPR, abuse,…
- Strong origin authentication appears to fail in the face of identity theft and end device capture
- Networks are not closed trust domains
- Is this the wrong question in the wrong place?

Common Constraints

Routing

- Routing systems operate within finite constraints
  - Some form of object abstraction is required to map the object domain into a smaller and more dynamically constrained routing domain
- Packet networks rely on per packet address lookups to determine local forwarding decisions
  - The abstraction is one of the imposition of hierarchies in the address plan where the hierarchy matches the physical topology

“One can route packets or politics, but probably not both”
(John Klensin)

Alternate Worlds?

- Is there anything else around?

  Not in the near term
- How long would a new design effort take?

  A decade or longer
- 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

So “extinction” is not very likely

What about “evolution”?

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

Some Technology Issues in IPv6

- The Address Plan
- Stateless auto-configuration
- Unique Local Addresses
- Flow Label
- QoS
- Security
- Mobility
- Multi-addressing
- Routing capabilities
- Revisiting endpoint identity and Network locator semantics

The Business Obstacles for IPv6

- Deployment by regulation or fiat has not worked in the past – repeatedly
- 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, address scarcity, and insecurity
- There is the prospect of further revenue erosion from simpler cheaper network models
- 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
- There is no current incremental revenue model to match incremental costs
- IPv6 promotion may have been too much too early – these days IPv6 may be seen as tired not wired
- Short term interests do not match long term common imperatives
- “Everything over HTTP” has proved far more viable than it should have

Meet the Enemy!

- “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

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

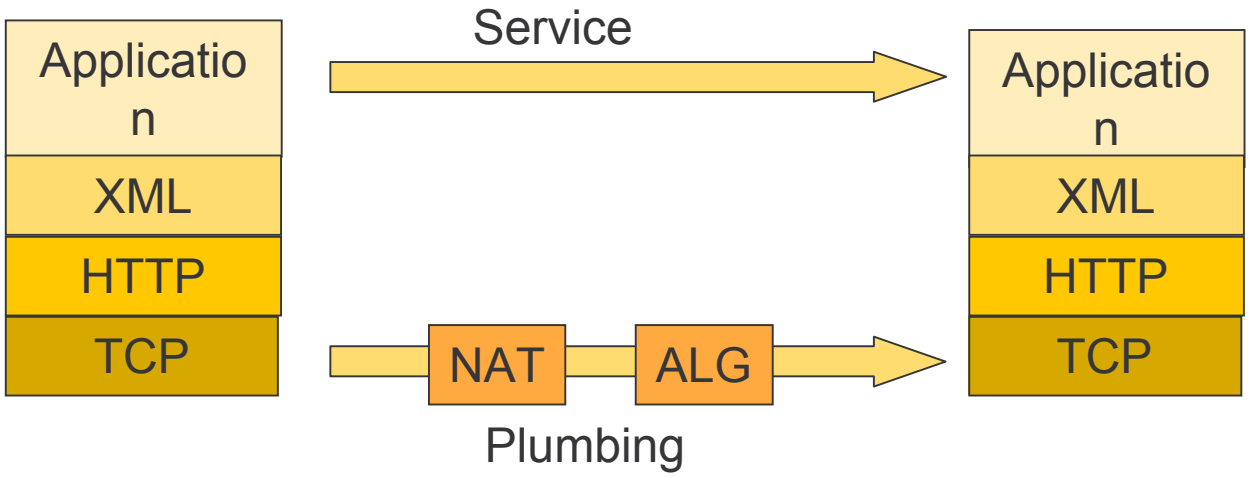

Maybe it's just business

- 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 and the costs of infrastructure deployment of IPv6

So “evolution” does not look that likely either

What about “revolution”?

Learning from IPv4

- IPv4 leveraged
  - cheaper switching technologies
  - more efficient network use
  - lower operational costs
  - structural cost transferral
- A compelling and revolutionary business case of cheaper services to consumers based on the PC revolution

IPv6?

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

- Only if we can further reduce IP service costs by a further 2 -3 orders of magnitude

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 Condition

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

- The silicon industry has made the shift from value to volume_years ago
- What will drive IPv6 is also a value to volume shift in the IP packet industry
  - The prospect of the V6 network embracing a world of trillions of chattering devices

IPv6 Revolutionary Leverage

- **Volume over Value**
  - Supporting a network infrastructure that can push down unit cost of packet transmission by orders of magnitude
  - V6 will 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 may well be “revolution” where IPv6 is forced into direct competition with existing IPv4+NAT networks

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

Maybe IPv6 is the catalyst towards shifting the Internet infrastructure industry a further giant leap into a future of commodity utility plumbing!

- Thank you!