The Case for IPv6

Extinction, Evolution or Revolution?

Geoff Huston
APNIC

A Long Time Ago .....

1991 - 1993 IETF ROAD effort to examine the two scaling issues of routing and addressing

There were three major outcomes of this exercise

- interim routing approach of provider-based address aggregation (**CIDR**)
- **IPv6** is defined using extended address headers in the packet
- Network Address Translators (**NATs**) are deployed in the Internet

IPv6 – the BGP view

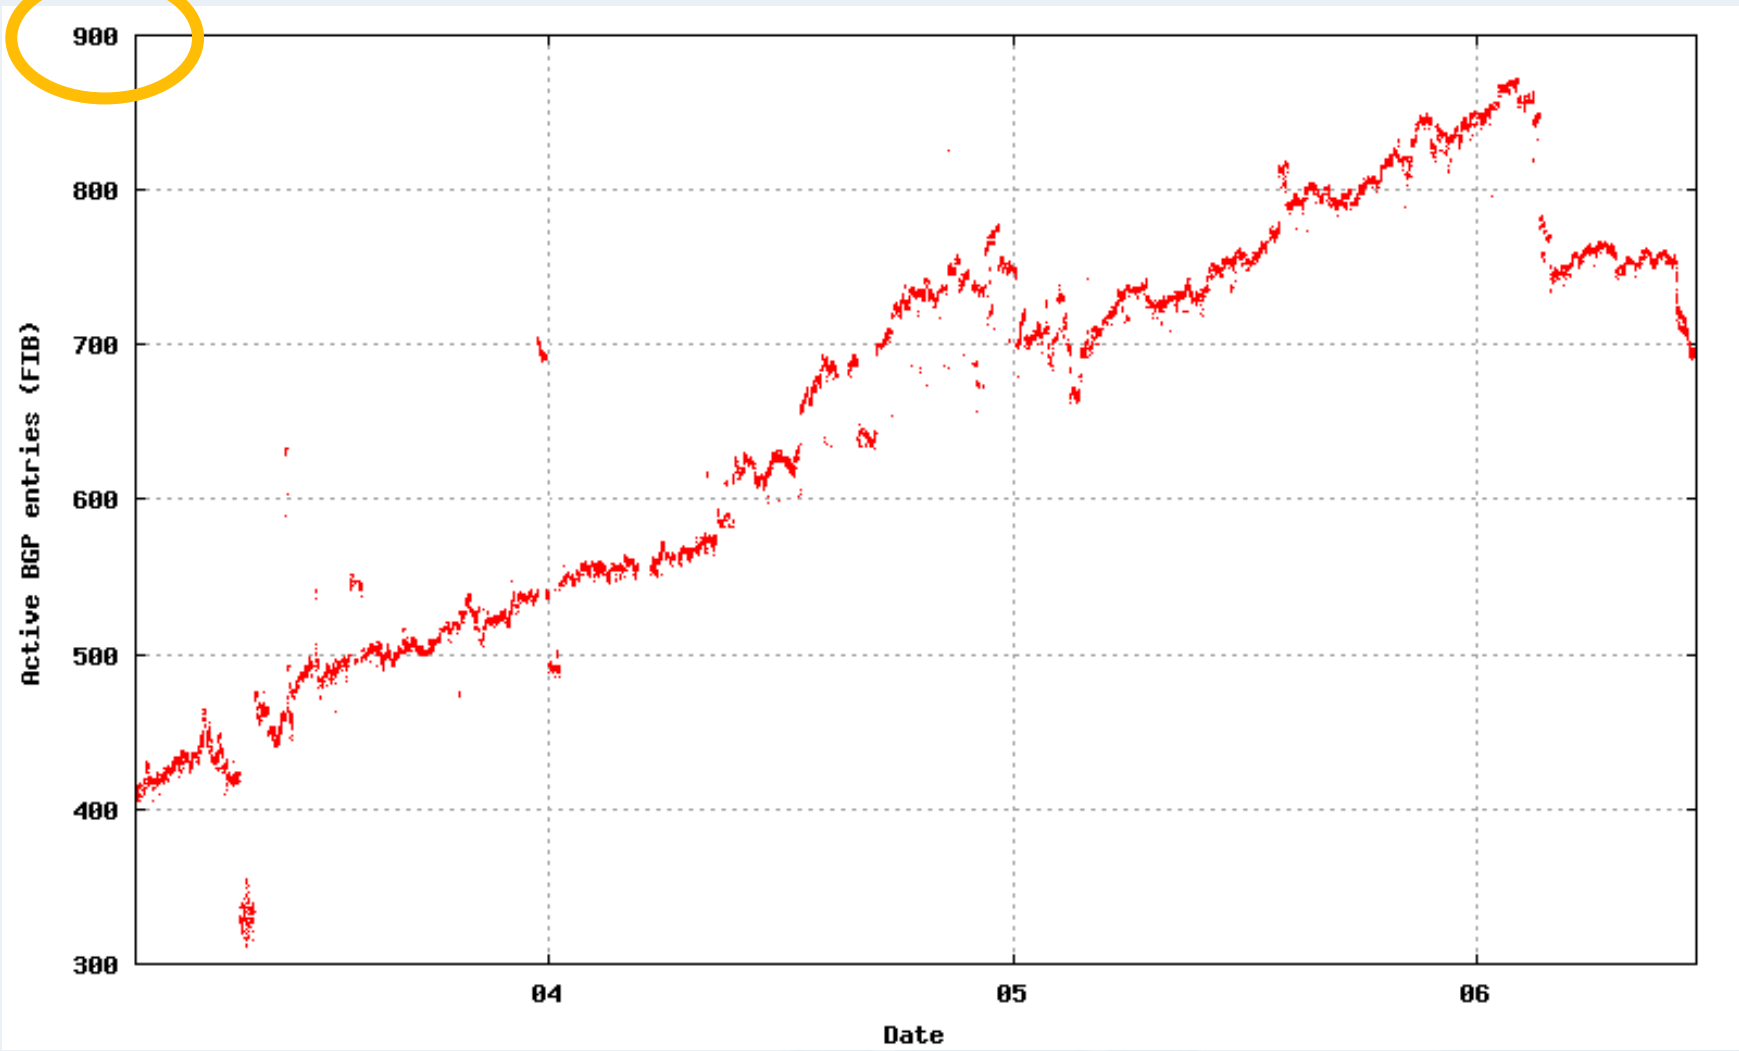

BGP Table Entries

IPv4

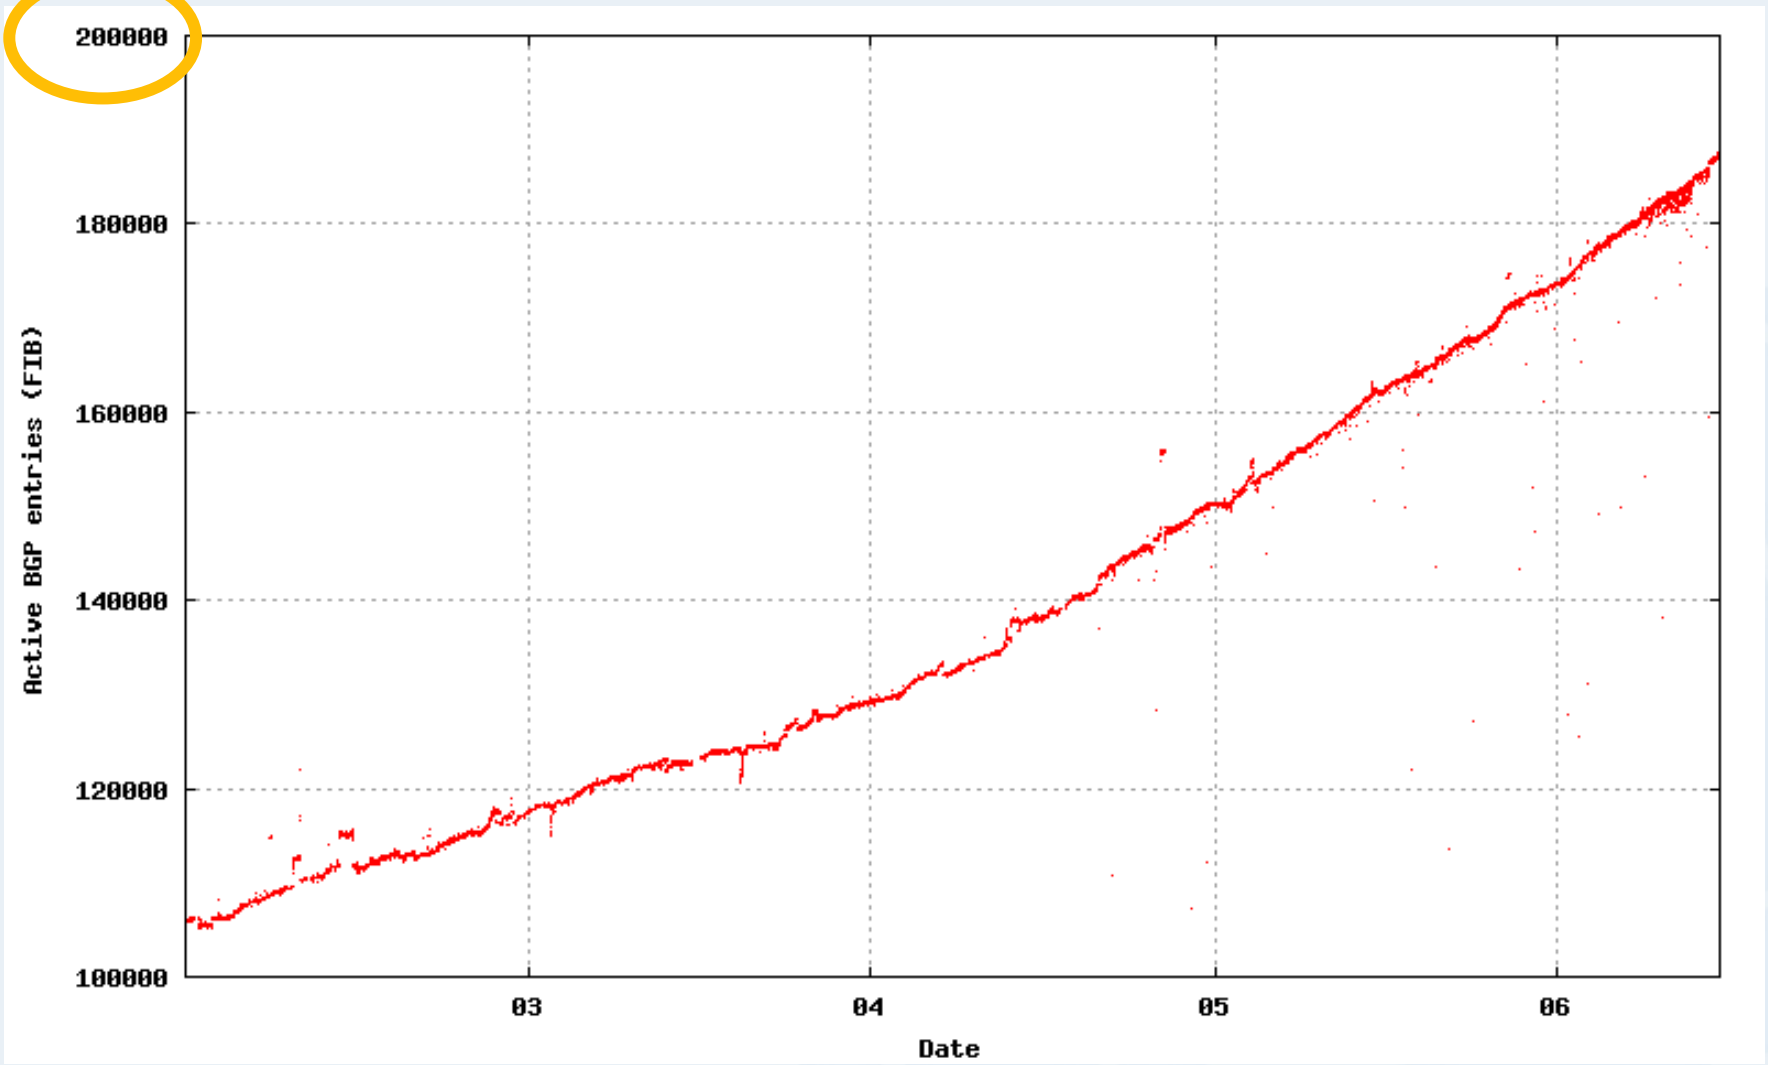

BGP Table Entries

IPv6 - AS Count

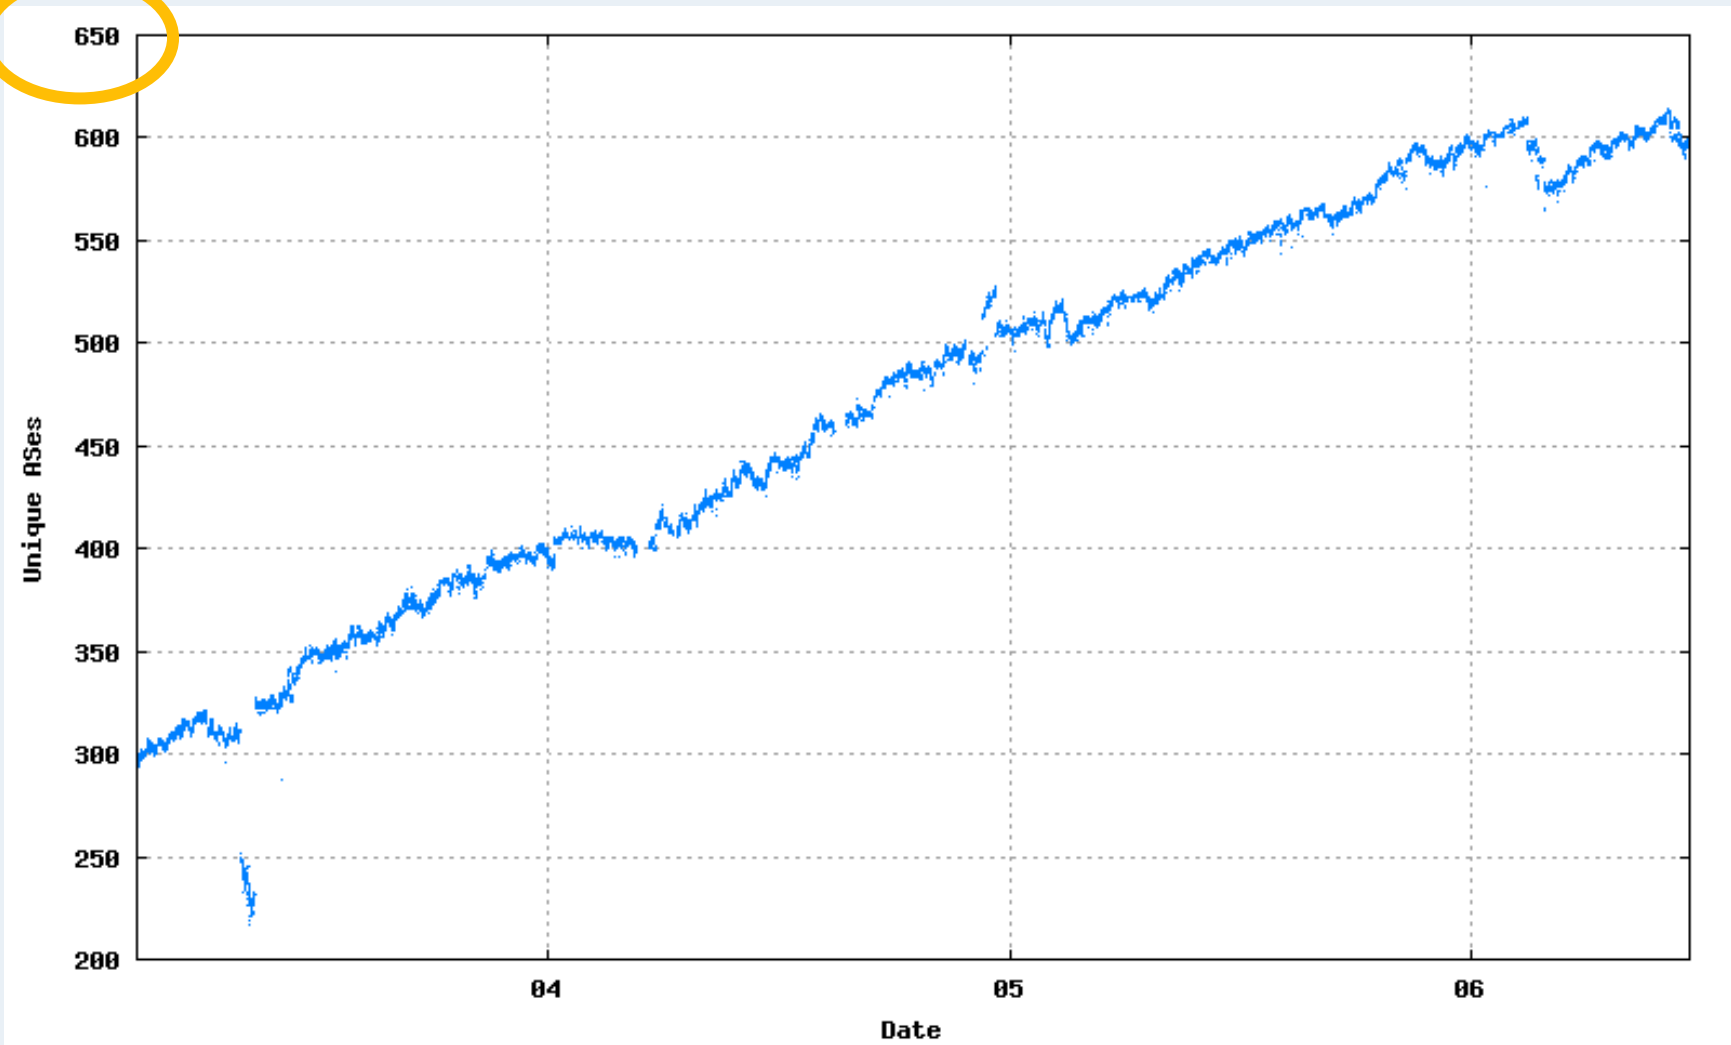

BGP AS Count

IPv4 - AS Count

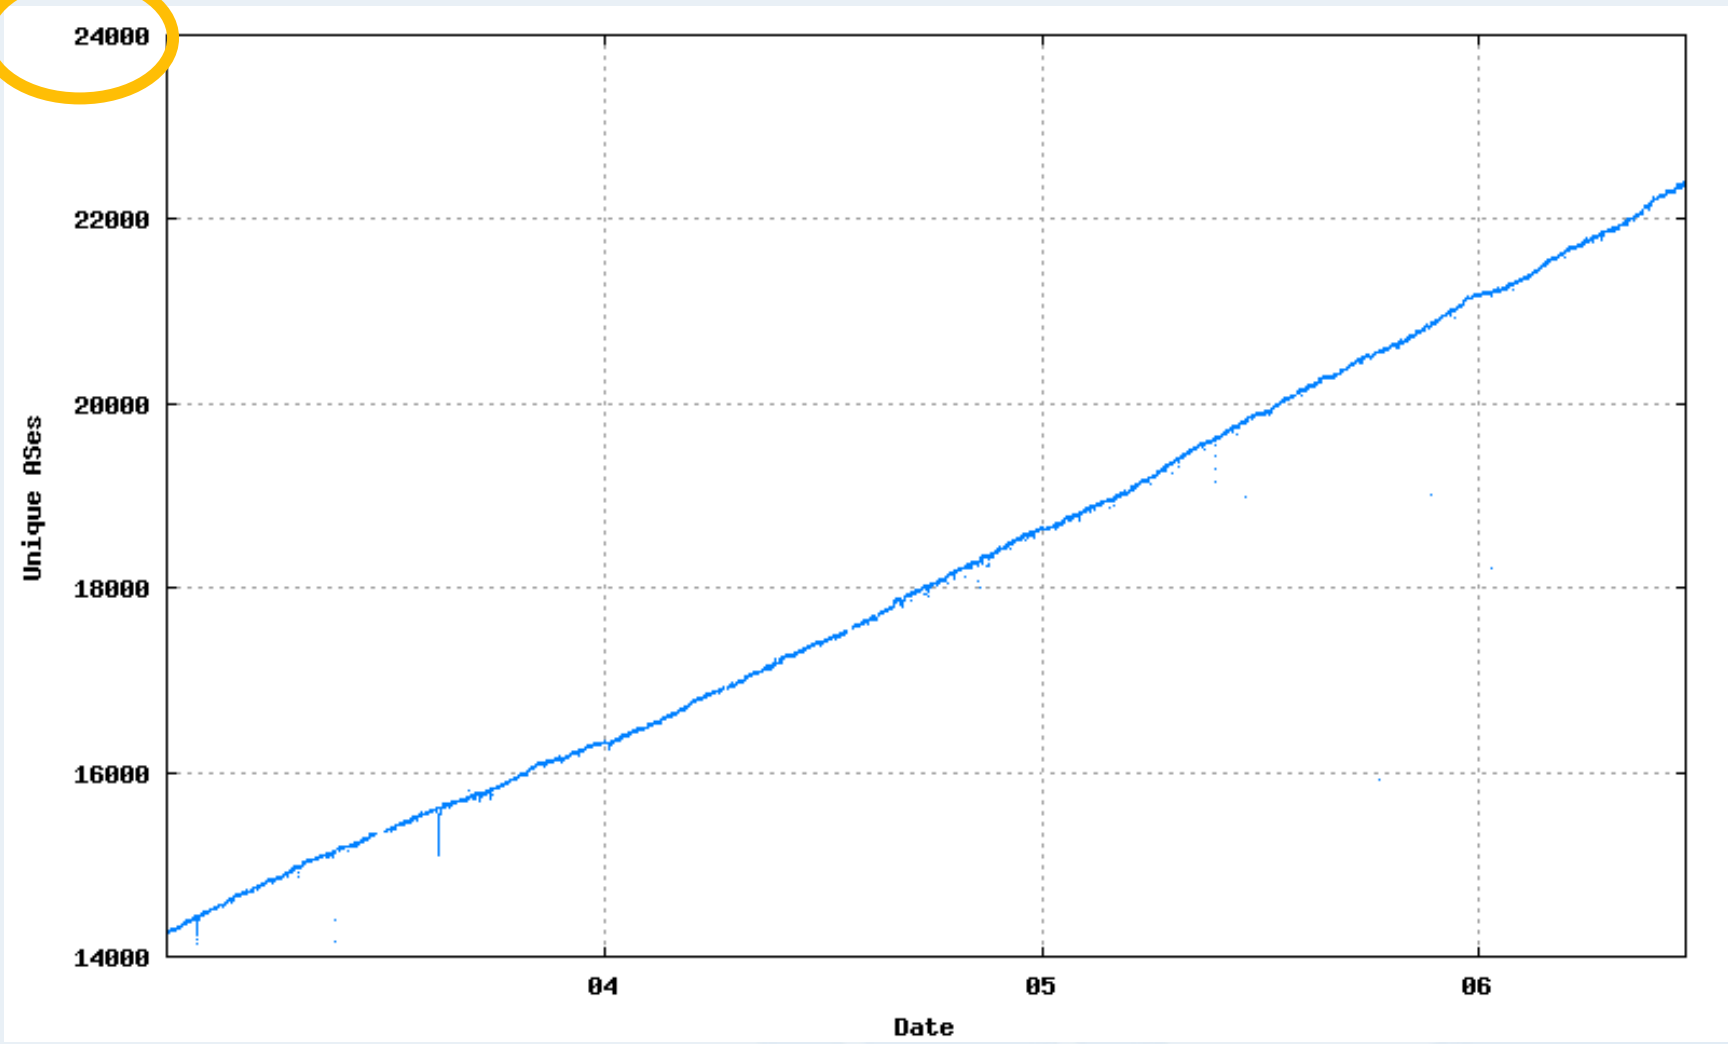

BGP AS Count

Is this a Problem?

- Available data indicates that IPv6 is not a significant component of today's global Internet industry
- So will we ever get to widespread IPv6 deployment? How?

IPv6 Adoption as an Evolutionary Process

- The Internet as an evolving lifeform or ecosystem
  - If IPv6 can offer clearly superior value propositions to the industry 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!

Is IPv6 adoption 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?

The Options for IPv6 Adoption

- Extinction?

The Case for IPv6 Extinction?

- The original IP architecture is dying – if not already terminally dead
  - Coherent transparent IP 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*, *turn*, *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
- IP was a neat idea, but the industry killed it!

Today

- We are engineering applications and services in an environment where NATs, firewalls and ALGs are assumed to be part of the IP plumbing
  - Client-initiated transactions
  - Application-layer identities
  - Agents and brokers to orchestrate multi-party rendezvous and NAT identification and traversal
  - Multi-party shared NAT state
- All this complexity just results in more fragile applications and lower revenue margins with higher risk

IPv6?

- We've all heard comments that:
  - IPv6 was rushed through the standards process
  - 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 *

(* same as V4 + NAT!)

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?

Alternate Worlds?

- Is there anything else around?
  - No - 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!)

So “extinction” is not very likely – there is simply no other option on our technology horizon

The Options for IPv6 Adoption

- ~~Extinction~~
- Evolution?

Should we evolve the Internet to use IPv6?

- 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
  - ISPs are not device manufacturers

- There appear to be no clear “early adopter” rewards for IPv6
  - Existing players have strong motivations to defer expenditure decisions -– because their share price is plummeting
  - New players have no compelling motivations to leap too far ahead of their seed capital
  - All players see no 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….

- We are not sending the right signals that this is ‘cooked and ready’
- We are still working on:
  - The IPv6 Address Plan
  - Aspects of IPv6 Stateless auto-configuration
  - Scoped Addresses
  - IPv6 Flow Label
  - IPv6 QoS
  - IPv6 Security
  - IPv6 Mobility
  - Multi-addressing in IPv6 hosts
  - Multi-homing in IPv6
  - IPv6 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
  - 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, address scarcity, and insecurity
  - And they are not going to go away quietly or quickly
- There is the prospect of further revenue erosion from simpler cheaper network models
  - Further share price erosion in an already gutted industry

The 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
- There is no current incremental revenue model to match incremental costs
  - Customers won't pay a higher service tariff for IPv6
- IPv6 promotion may have been too much too early – these days IPv6 may be seen as tired not wired
  - Too much powerpoint animations?
- Short term individual interests do not match long term common imperatives
  - The market response is not driven by longer term concerns
- "Everything over HTTP" has proved far more viable than it should have

More Obstacles

- “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

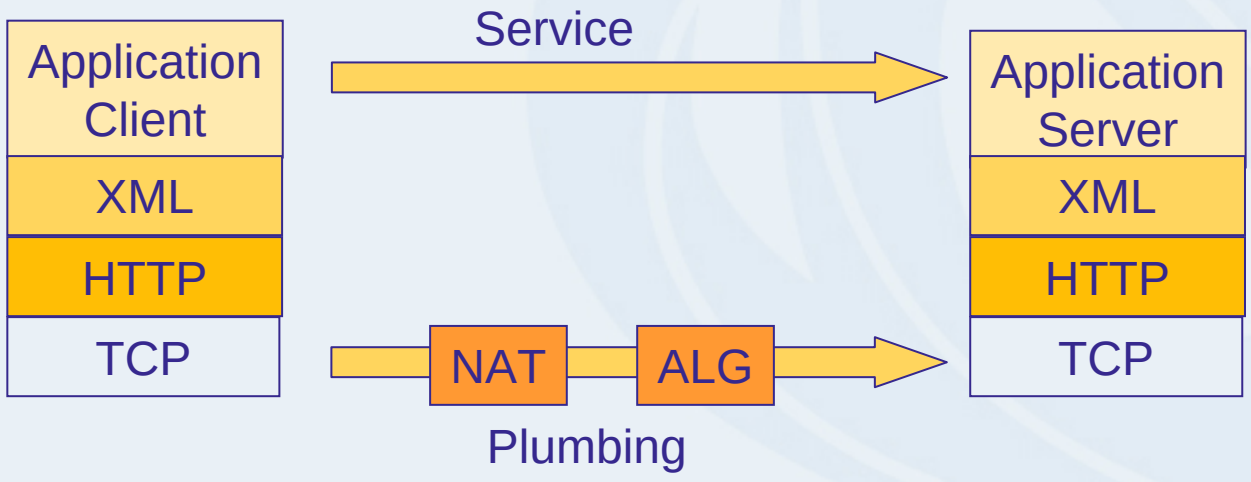

Maybe it's just deregulation and market behaviours?

- 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

So “evolution” does not look that likely either

The Options for IPv6 Adoption

- ~~Extinction~~
- ~~Evolution~~
- Revolution?

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

IPv6?

- IPv6 represents an opportunity to embrace the communications requirements of a device-dense world
  - Way much more than PCs
  - Device population that is at least some 2 – 3 orders of magnitude larger than today's Internet (100 to 1,000 times larger than today)

- 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 may well be “revolution” where IPv6 is forced into direct competition with existing IPv4 networks

- And the primary leverage here is one of *cheaper* and *bigger* network infrastructure, and not necessarily “better” or “smarter” networks

If all IPv6 can offer is just IPv4 with bigger header fields then the ‘transition’ into IPv6 has already stalled and its unclear how it will ever regain industry momentum

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

Maybe IPv6 is the catalyst for a future of commodity utility plumbing in a silicon dense world

Thank You