The Evolution of the Internet and IPv6

Geoff Huston
APNIC
Australian IPv6 Summit
31 October 2005

IPv6 - the BGP view since 2003

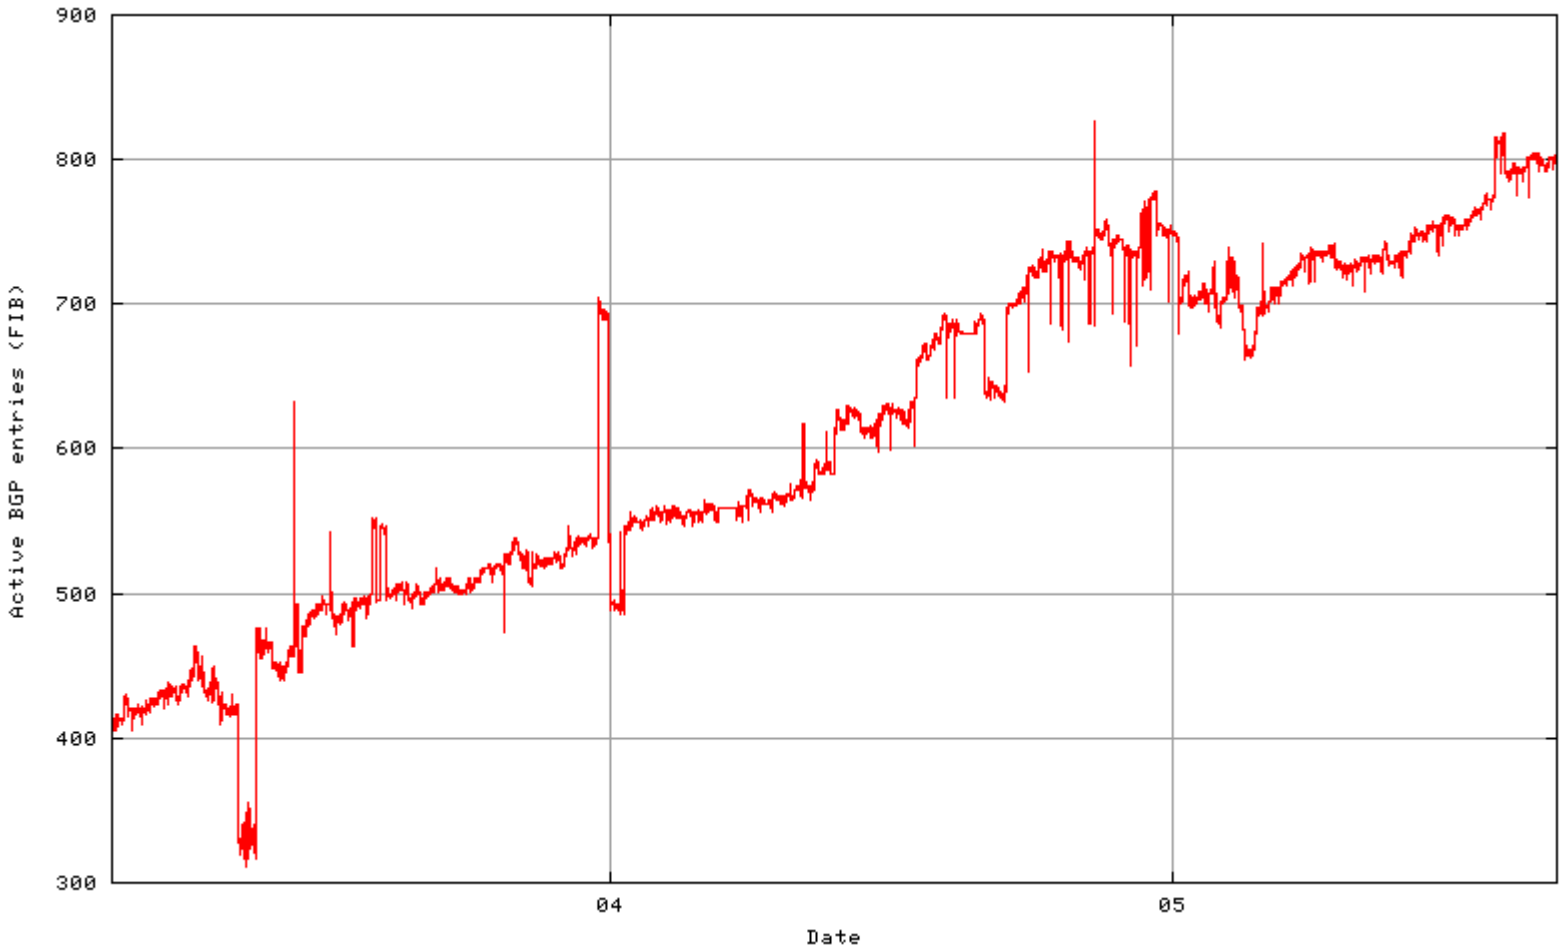

IPv4 – the BGP view since 2003

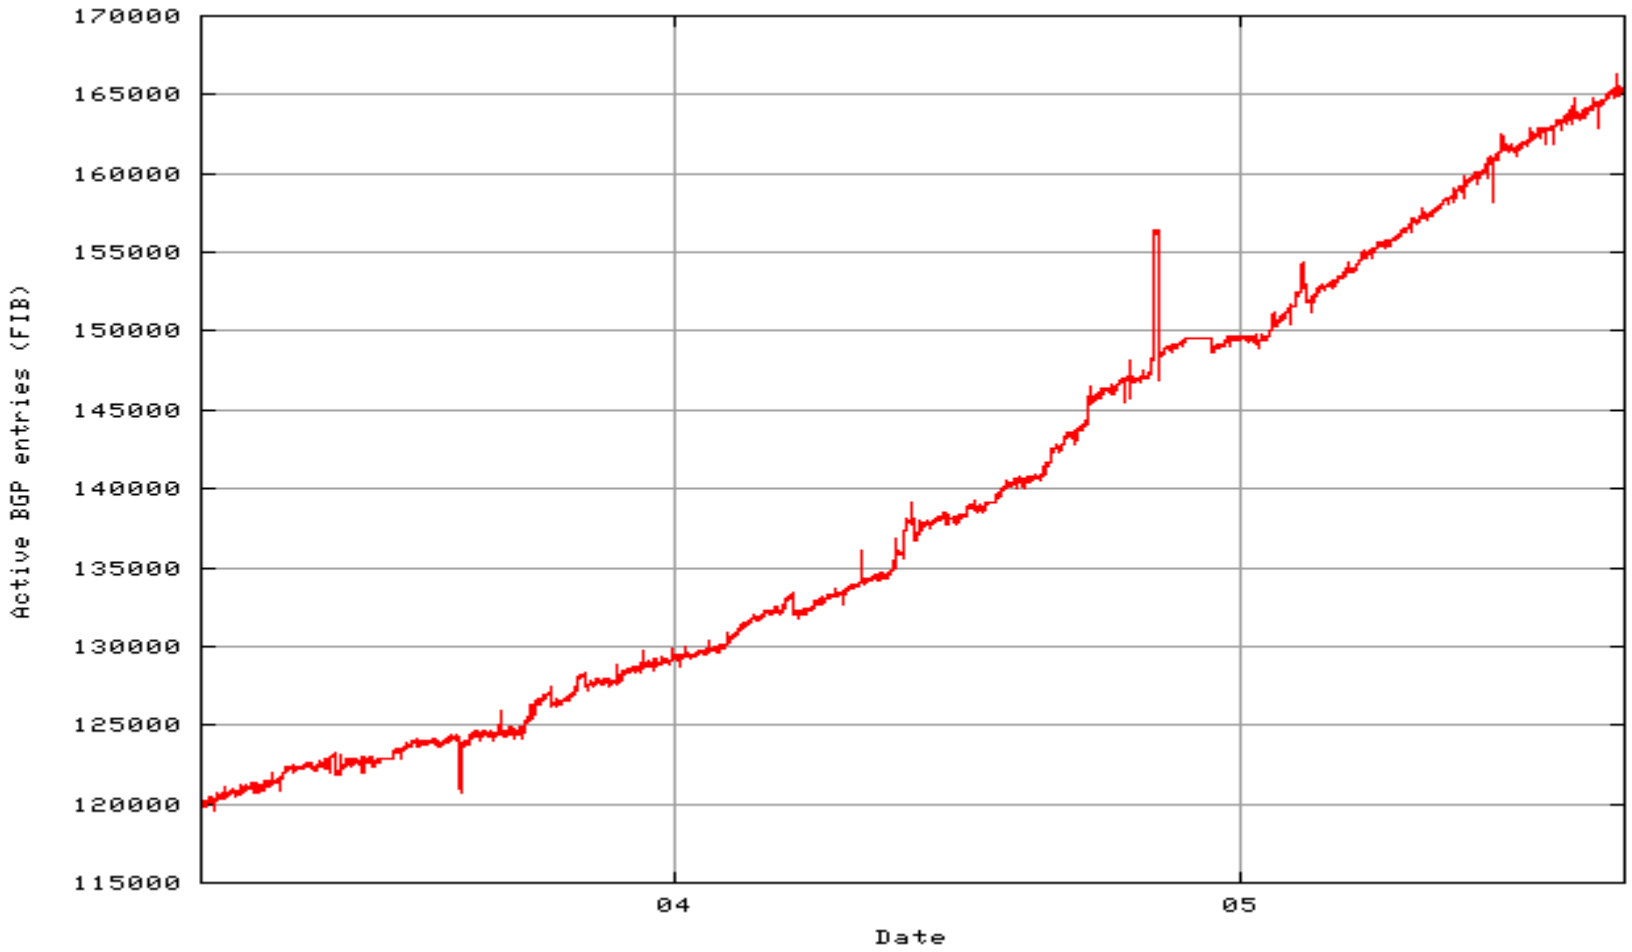

IPv6 Adoption – AS Count

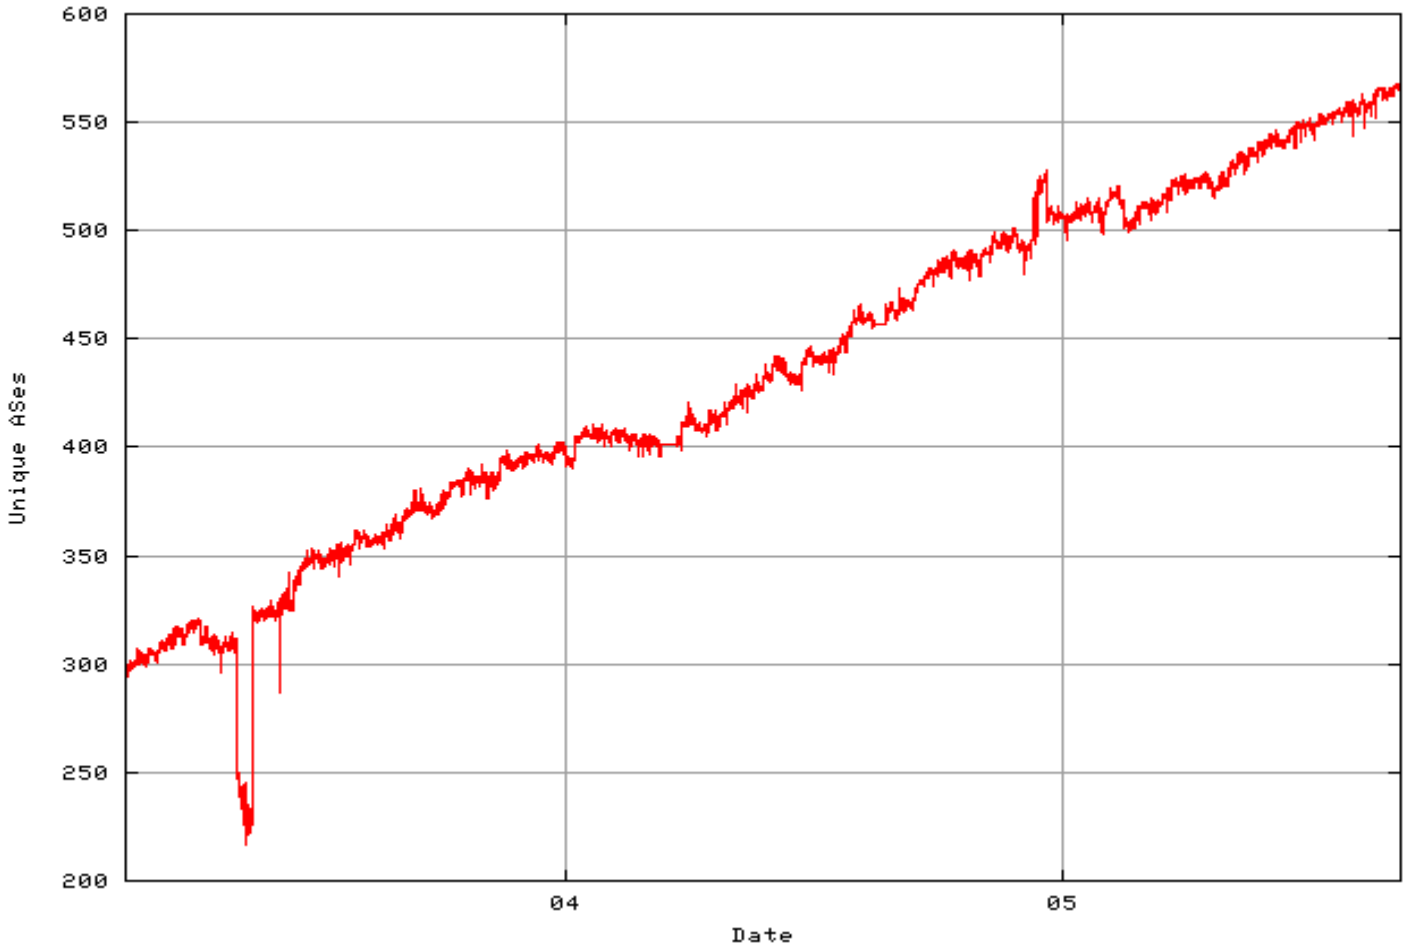

IPv4 Expansion – AS Count

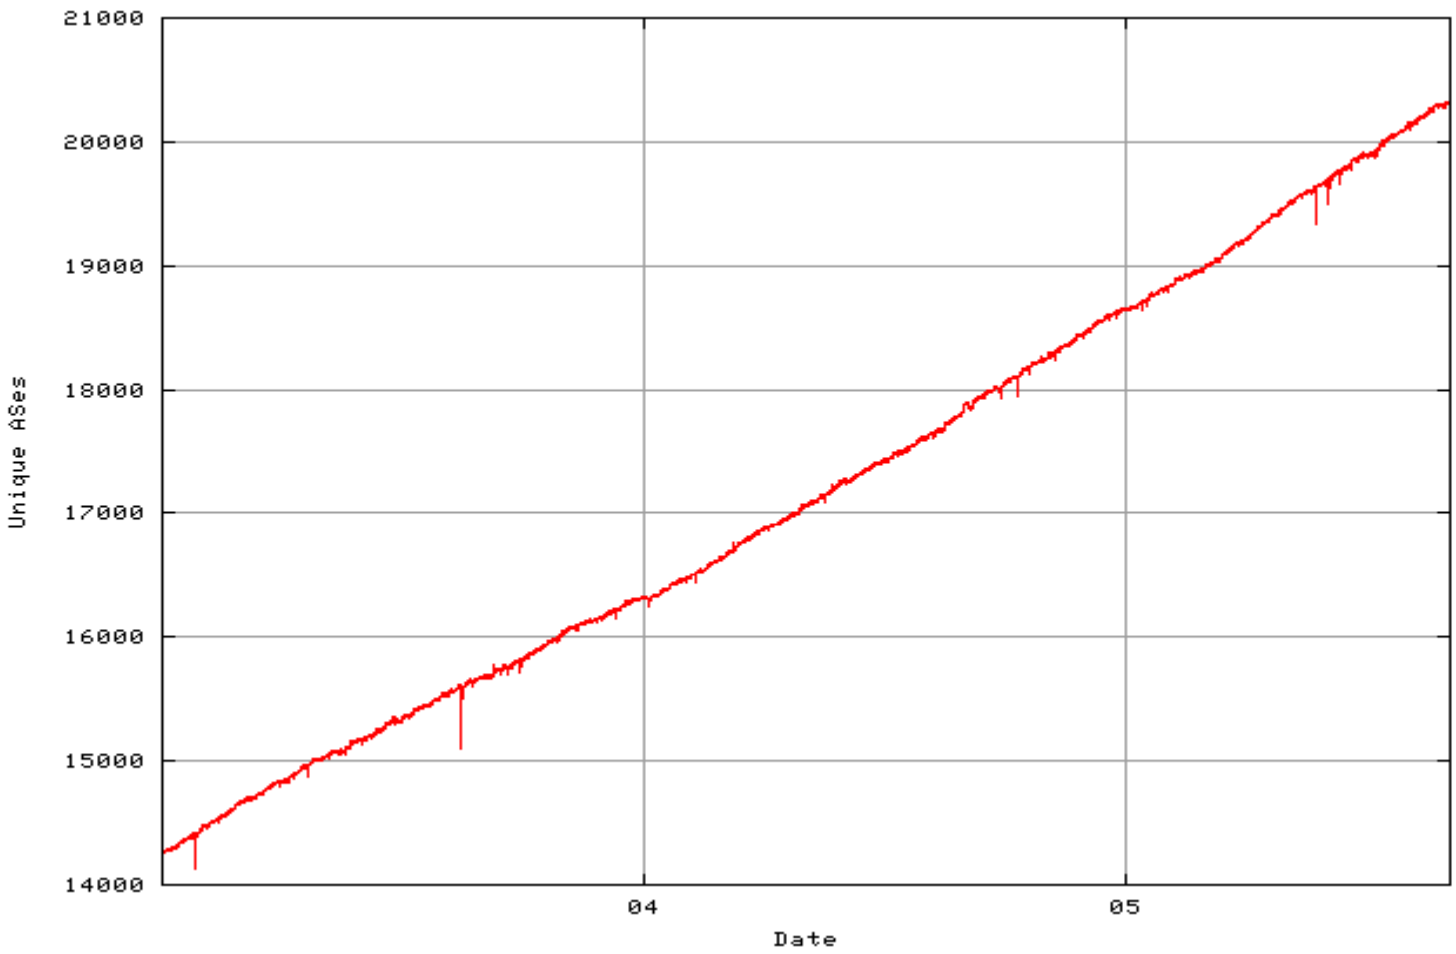

IPv6 vs IPv4 Rates – AS Count

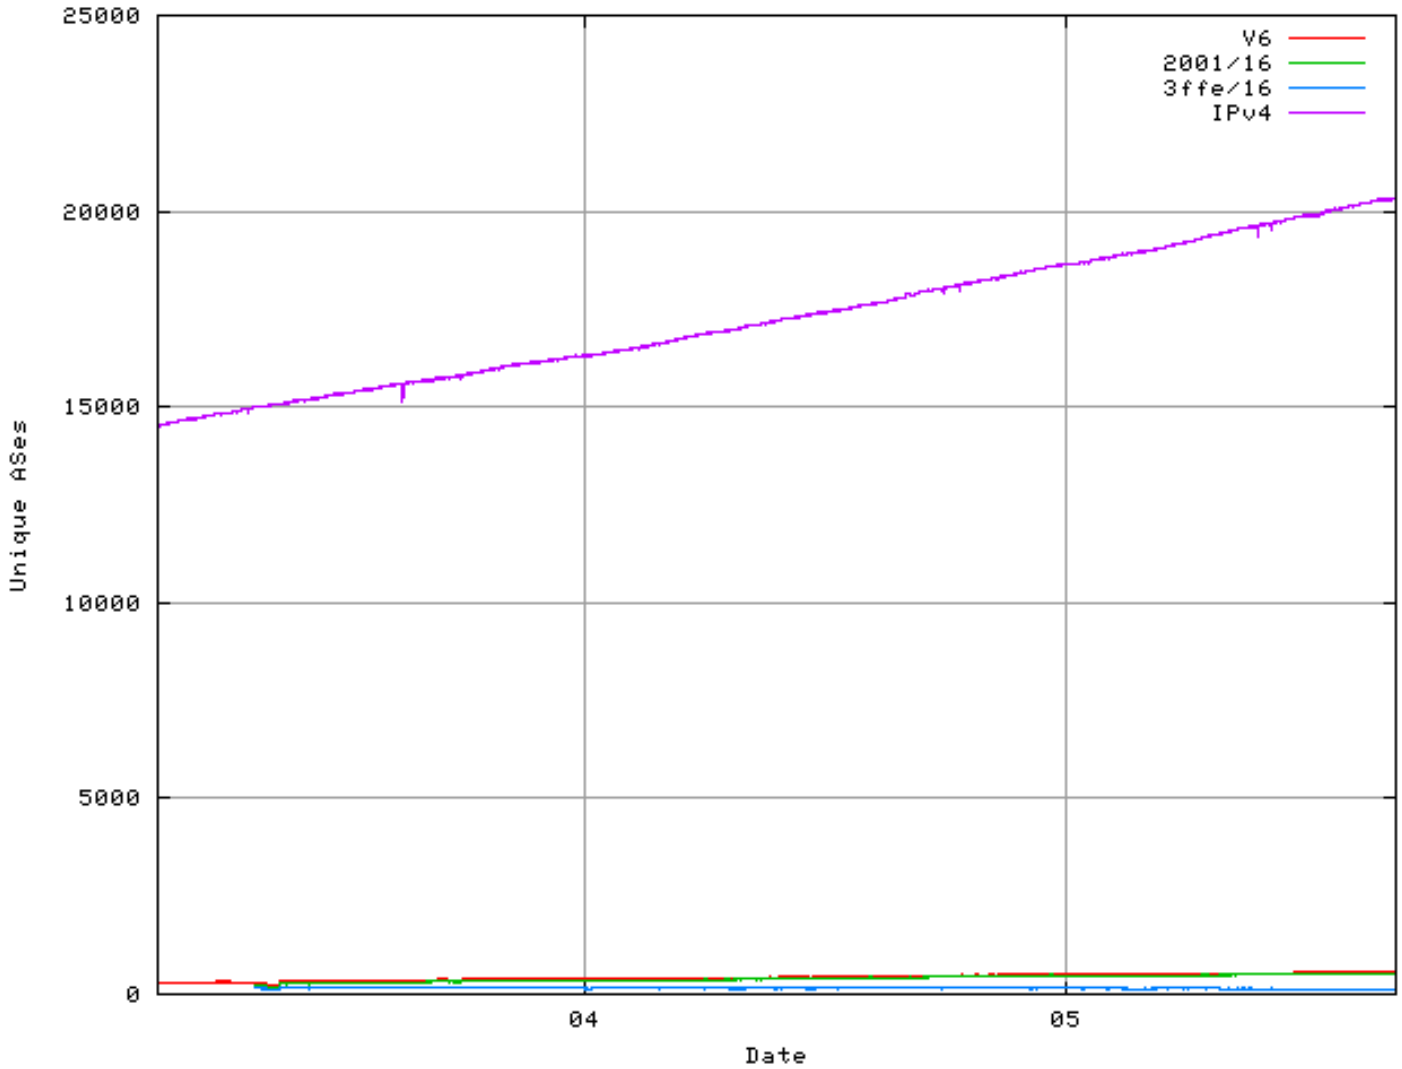

Innovation and Conservatism

- We’ve learned that optimism is no substitute for knowledge and capability within this industry
- But without optimism, innovation is stifled

- Current conservative period of consolidation rather than innovative expansion
  - 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
  - Accompanied by greater emphasis of financial returns from existing infrastructure investments

Is IPv6 as an innovation OBE?

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?

What is the story with IPv4?

- 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!

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 almost as much as megabit DSL access

We are running out of low cost 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. i.e. the price reflects the relative scarcity

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 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 higher operational margins

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 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 -– 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

IPv6?

- We’ve all heard views 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?

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 whole ‘control’ thing in network architecture just 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 a rich 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 approximately matches the physical topology

"*One can route packets or politics, but probably not both*"
John Klensin

"*We can't route money*"
Dave Clark

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!)

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

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

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
  - Unique Local Addresses (whatever they may be today!)
  - Flow Label
  - QoS
  - Security
  - Mobility
  - Multi-addressing
  - Multi-homing
  - 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

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!
- 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 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 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

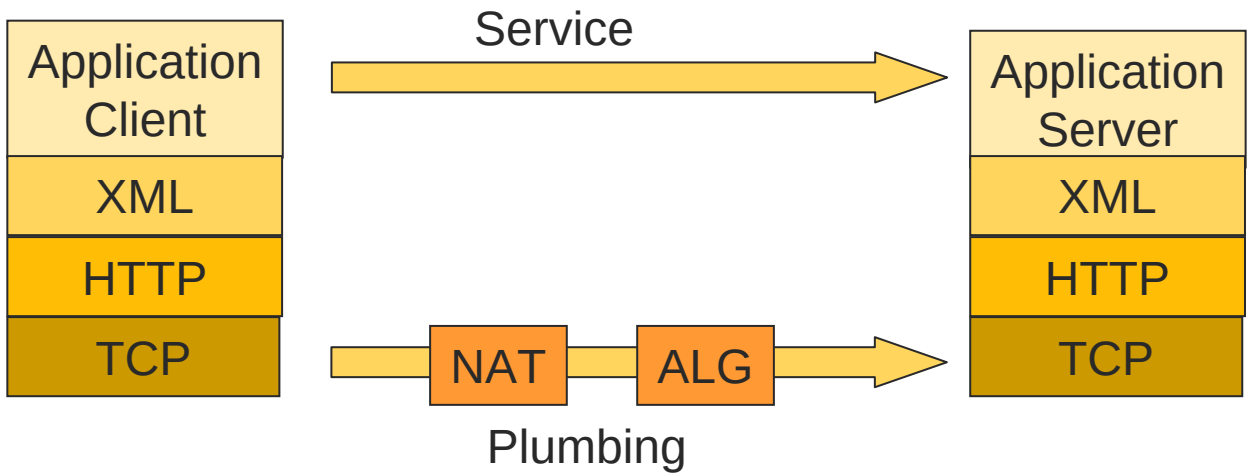

Maybe it's just deregulation

- 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

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

- IPv4 represented a compelling and revolutionary business case of stunningly cheaper and better services to end consumers, based on the silicon revolution

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

- 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 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

- The silicon industry has made the shift from **value** to **volume** years ago
- What will drive IPv6 deployment in a device rich world is also a radical and **revolutionary** value to volume shift in the IP packet carriage industry

IPv6 Revolutionary Leverage

- **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

But it won't be easy

capital distribution problem

(the ones who need to innovate in the core don't have capital)

INNOVATOR	EPS ($)	MKT CAP ($B)
MCIW	-11.22	6.5
SPRNT/NXTL	-0.31	34
VERIO/NTT	1.98	71.6
LEVEL3	-0.74	1.9
SBC/T	1.41	78
QWEST	-0.45	7.7
COGENT	-7.42	0.2
GLBC	-13.84	0.3
SAVVIS	-0.90	0.12
ABOVENET	n/a	n/a
WILTEL	n/a	n/a
TELEGLOBE	-0.74	0.2
C&W	0.70	4.7B
TWTELCOM	-1.12	1.0
(TWARNER)	0.48	82
XO	-2.18	0.4

INNOVATOR	EPS ($)	MKT CAP ($B)
CISCO	0.87	108
GOOGLE	3.41	97
AMAZON	1.25	19
YAHOO	1.07	49
EBAY	0.73	51
JUNIPER	0.53	13
APPLE	1.56	47.
INTEL	1.33	141
VERISIGN	0.93	6.15
DELL	1.27	76.3
MICROSOFT	1.12	269B

source: finance.yahoo.com, 25 oct 2005

- So it looks like the IPv6 future may well 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 IPv6 is the catalyst towards shifting the Internet infrastructure industry a further giant leap into a future of commodity utility plumbing!

And while you many not have a happy shareholder who is still expecting $5.25 a share and may have to live with something much much lower, at least you have some form of a future - as against none whatsoever!

- Thank you