Beyond the IPv4 Internet

Geoff Huston

Chief Scientist, APNIC

The IETF's ROAD Trip

▸ By 1990 it was evident that IPv4 was not going to have a large enough address span for long term deployment

▸ And the routing architecture was not able to scale indefinitely

▸ The combined ROuting and Addressing effort took up much of the IETF's attention in the period 1991 – 1994

▸ There were a number of outcomes – some intentional, some accidental

ROAD Outcomes

▸ **Short Term mitigation**

  ▸ Drop address classes from the address plan to decrease address consumption rates

  ▸ Adopt provider-based addressing to increased routing aggregation

▸ **Longer Term approach**

  ▸ Extend the address size in IP by a factor of 4

▸ **Accidental Outcome**

  ▸ NATs

ROAD Outcomes

▸ **Short Term mitigation**

  ▸ Drop address classes from the address plan to decrease address consumption rates *DONE!*

  ▸ Adopt provider-based addressing to increased routing aggregation

▸ **Longer Term approach**

  ▸ Extend the address size in IP by a factor of 4

▸ **Accidental Outcome**

  ▸ NATs

ROAD Outcomes

▶ **Short Term mitigation**

  ▶ Drop address cl___s from the address plan to decrease address co_____on rates ~~DONE!~~

  ▶ Adopt provi_____d addressing to increased routing aggregatio_ ~~DONE!~~

▶ **Longer Term approach**

  ▶ Extend the address size in IP by a factor of 4

▶ **Accidental Outcome**

  ▶ NATs

ROAD Outcomes

- **Short Term mitigation**
  - Drop address classes from the address plan to decrease address consumption rates ~~DONE!~~
  - Adopt provider based addressing to increased routing aggregation ~~DONE!~~

- **Longer Term approach**
  - Extend the address size in IP by a factor of 4

- **Accidental Outcome**
  - NAT (over) DONE!

ROAD Outcomes

- Short Term mitigation
  - Drop address ~~classes~~ from the address plan to decrease address co~~nsumption~~ rates _DONE!_
  - Adopt provide~~r base~~d addressing to increased routing aggregation _DONE!_
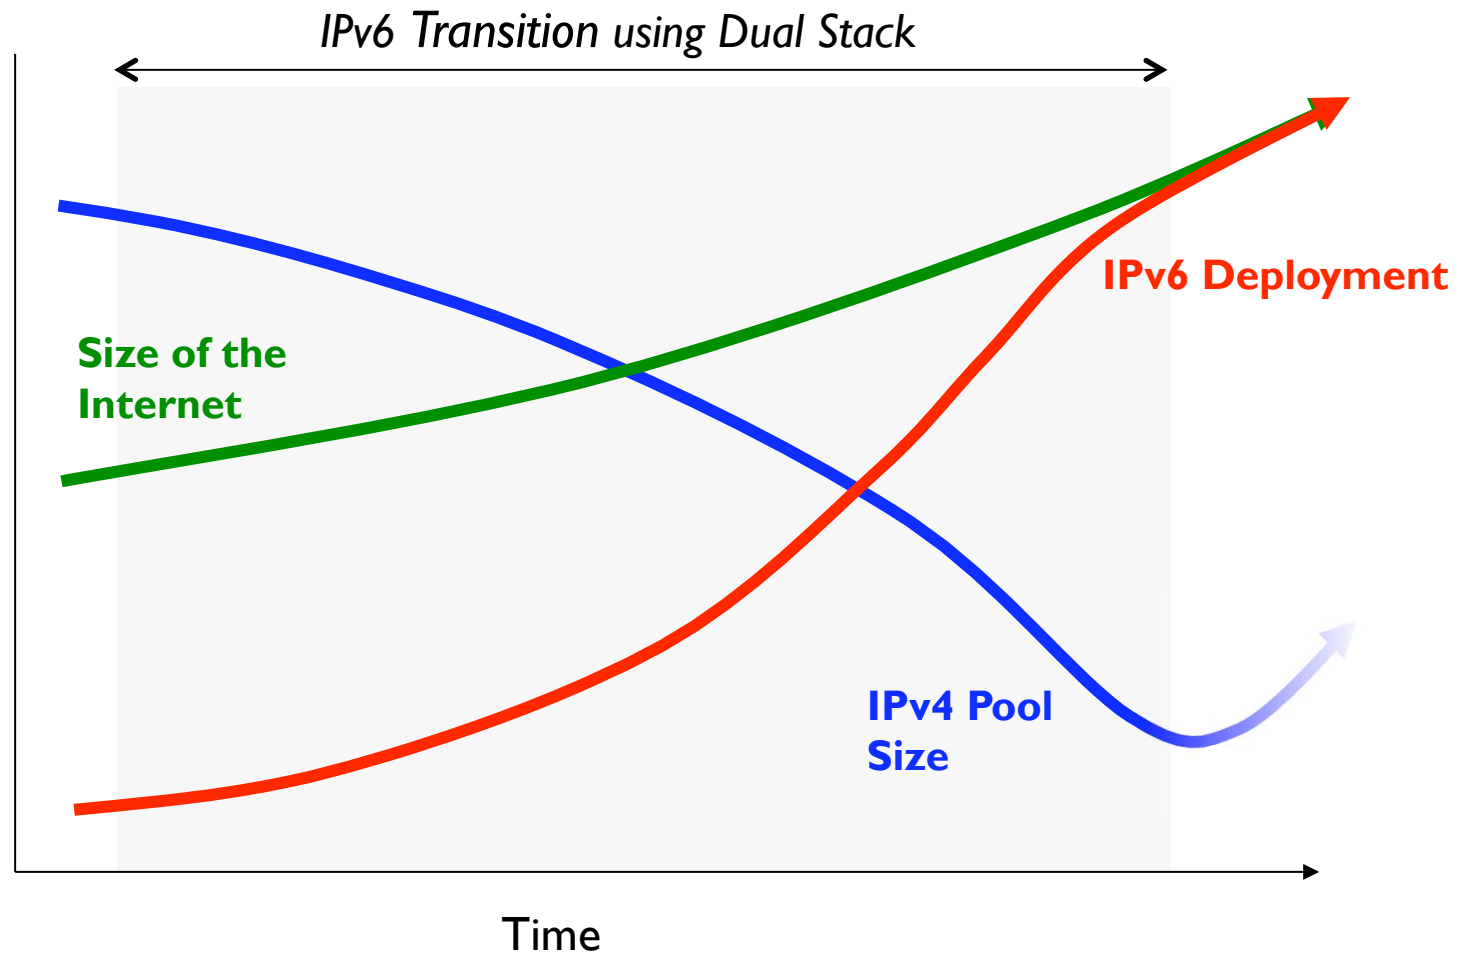- Longer Term approach
  - Extend t~~he IP address pool~~ by a factor of 4 _(Transition to IPv6)_
- Accidental Outcomes
  - NAT _(over) DONE!_

The Original Plan for IPv6 Transition

IPv6 Transition using Dual Stack

Size of the Internet

IPv6 Deployment
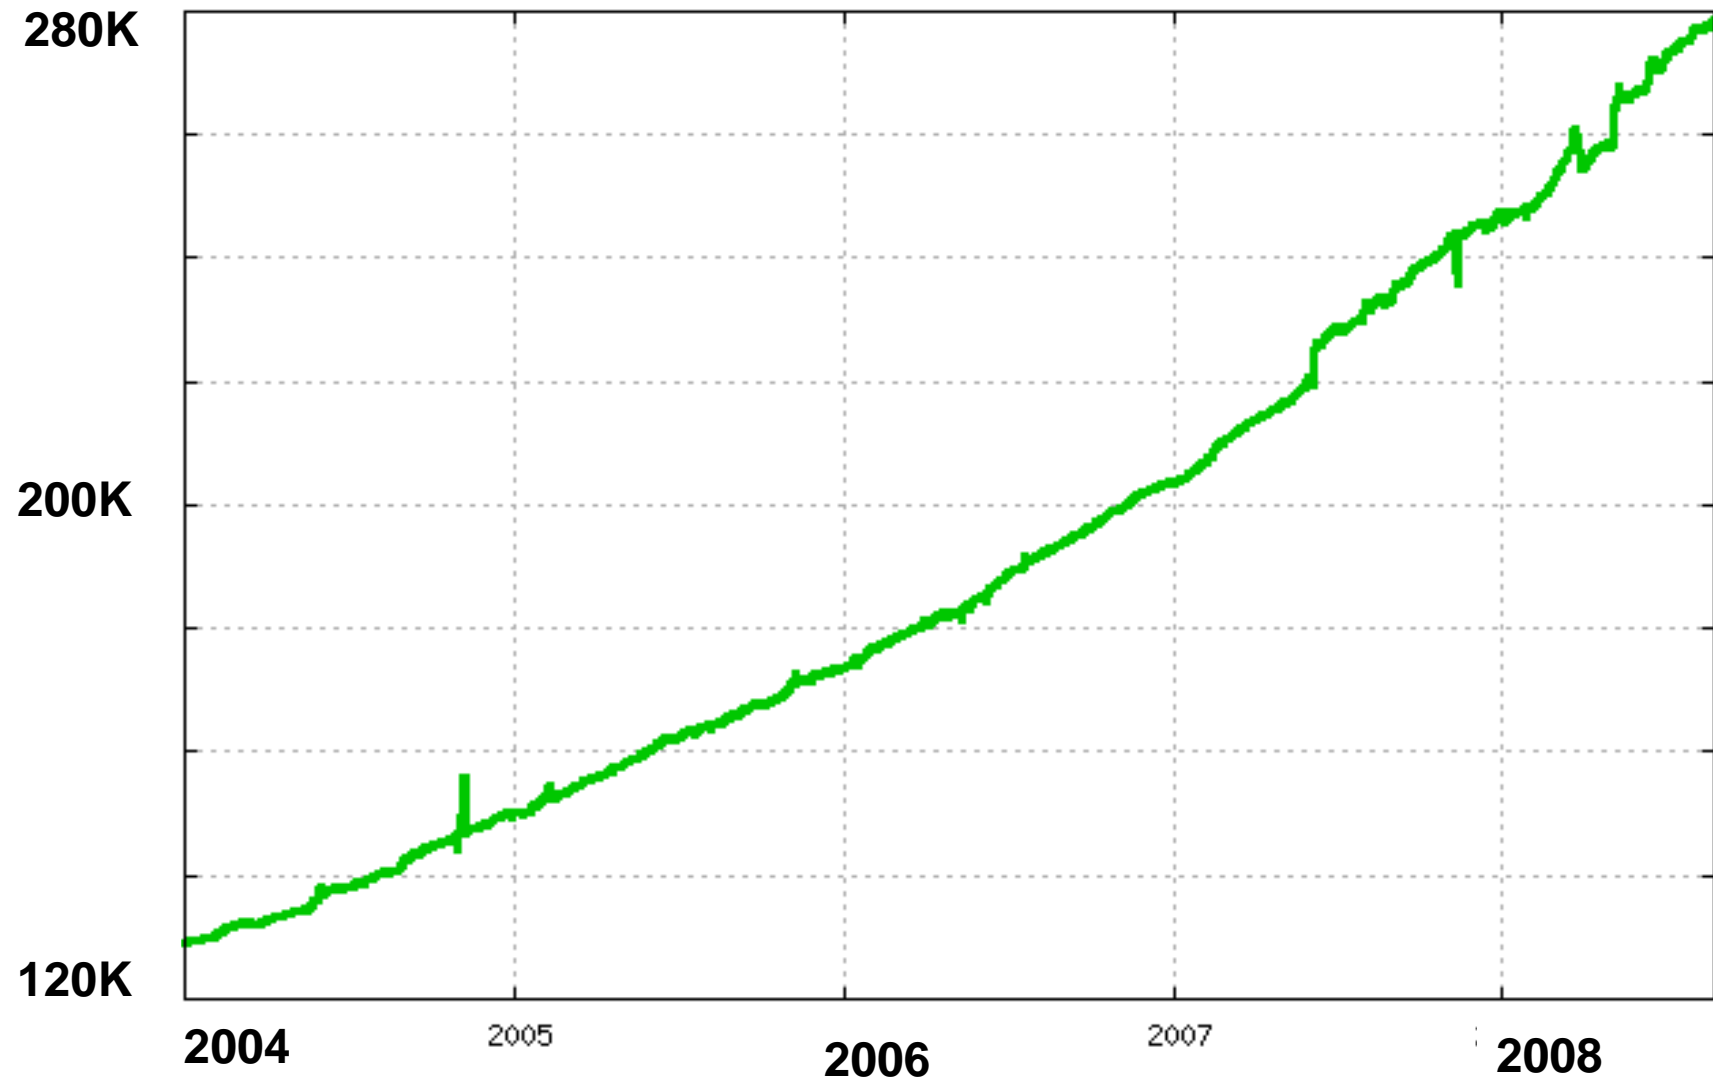
IPv4 Pool Size

Time

How are we doing in this plan?

▸ **Can we provide some measurements about where we are with IPv6 deployment across the entire Internet?**

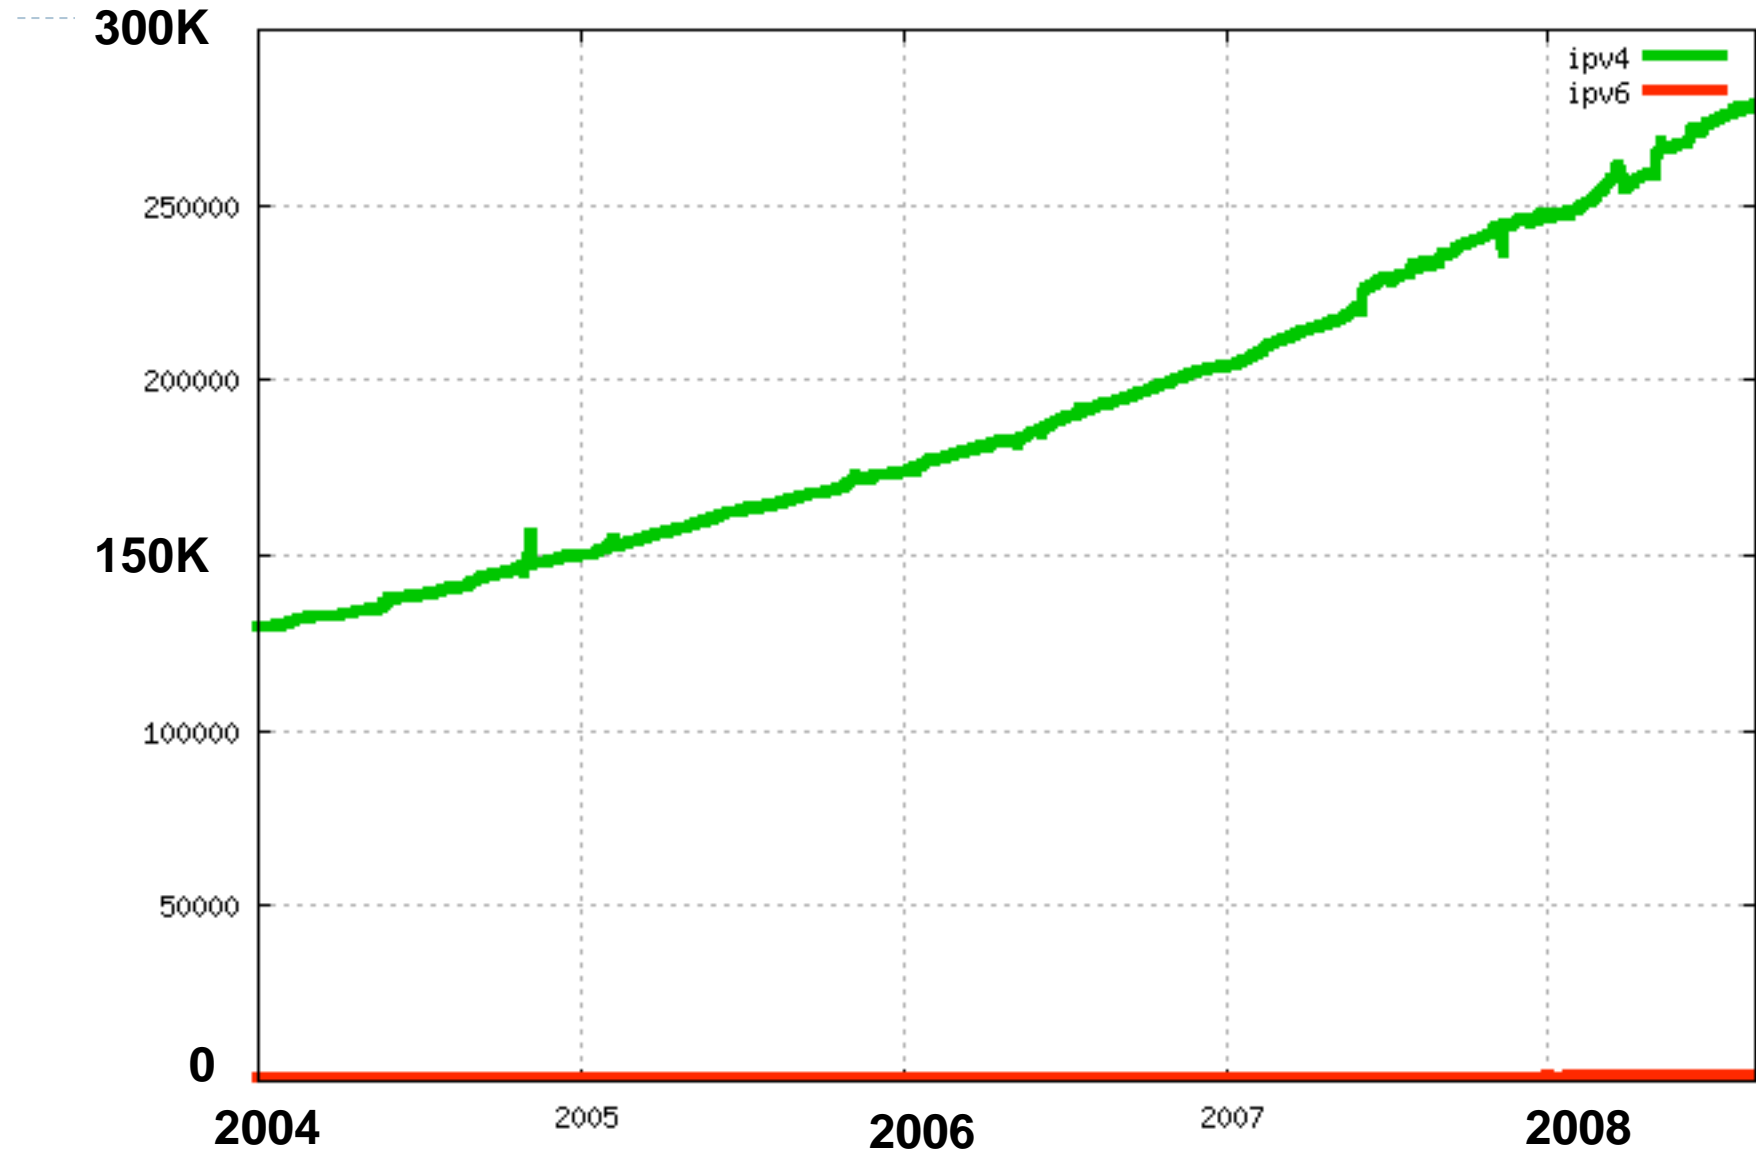 ▸ What measurements are useful?

  ▸ What data sets are available?

Routing Measurements:
The BGP view of IPv6

The BGP view of IPv4

BGP: IPv6 and IPv4

BGP IPv6 : IPv4

BGP IPv6 : IPv4

Some Observations and Measurements
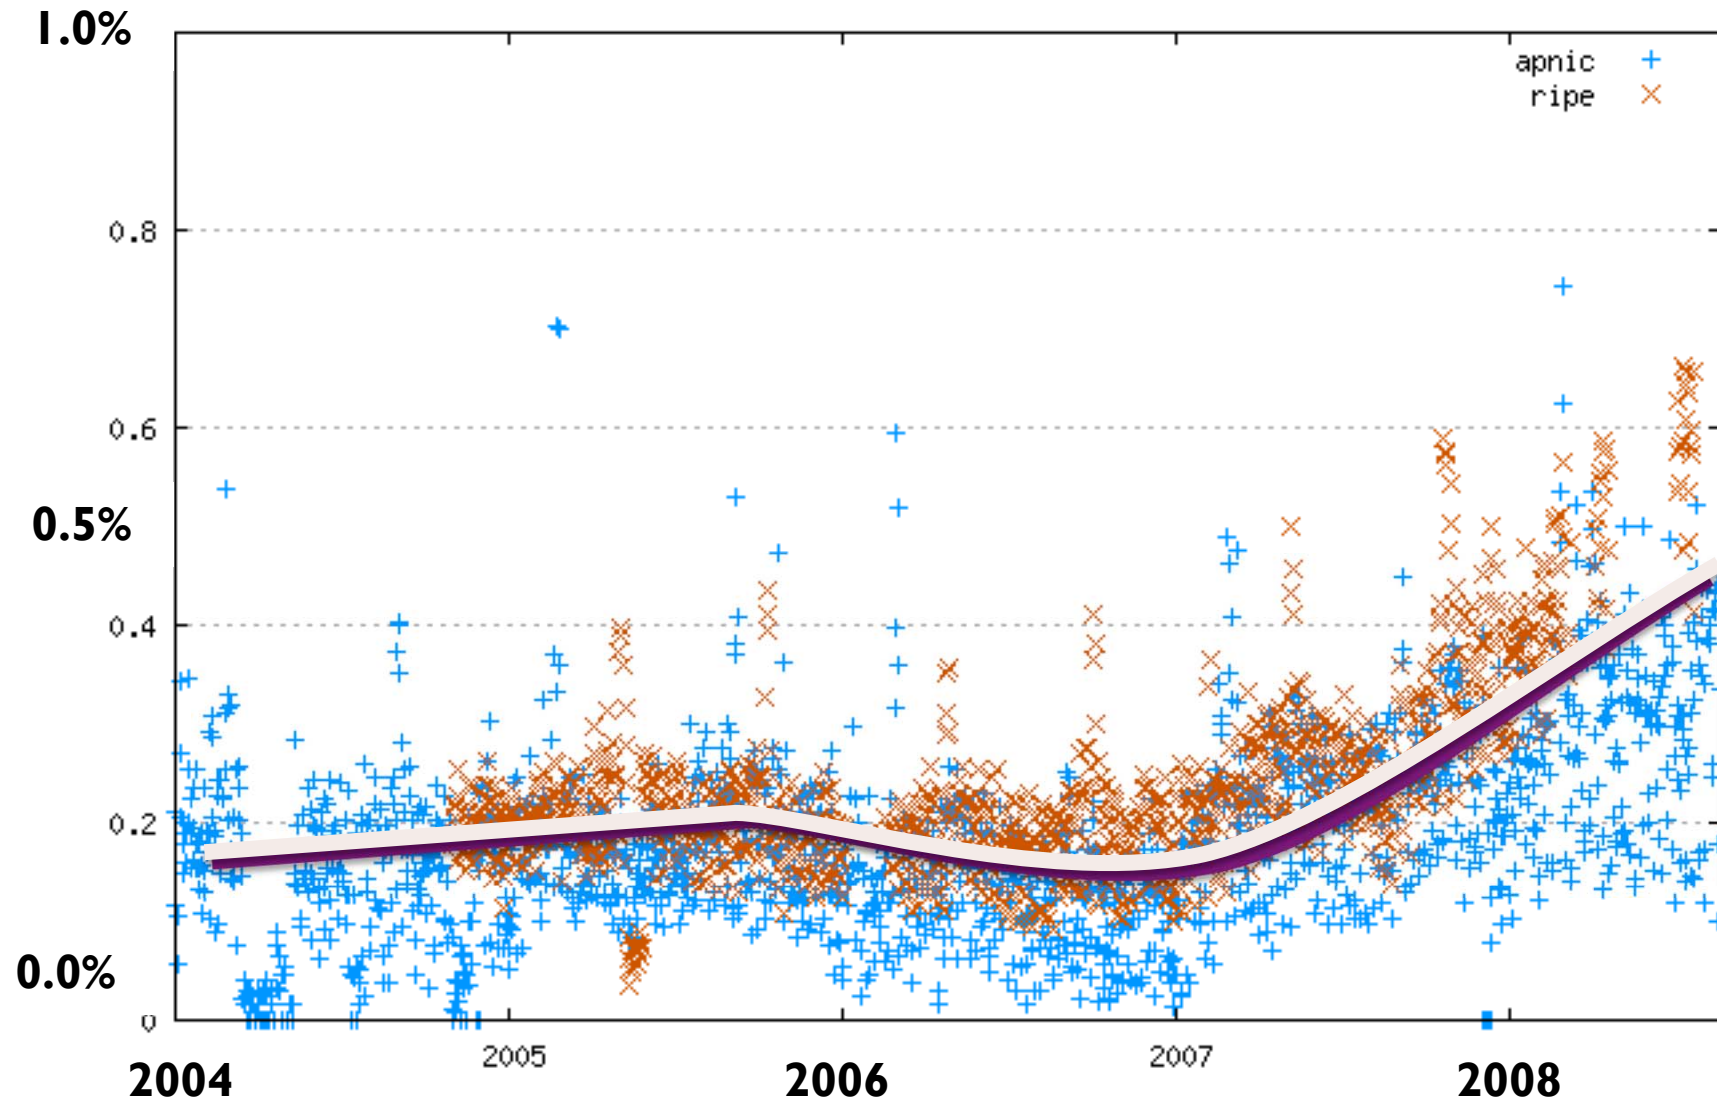
▸ IPv6 represents 0.5% of all BGP routes

Some Observations and Measurements

▸ IPv6 represents 0.5% of all BGP routes

But how significant is a relative growth in IPv6 routing entries from 0.4% to 0.5% over 18 months?
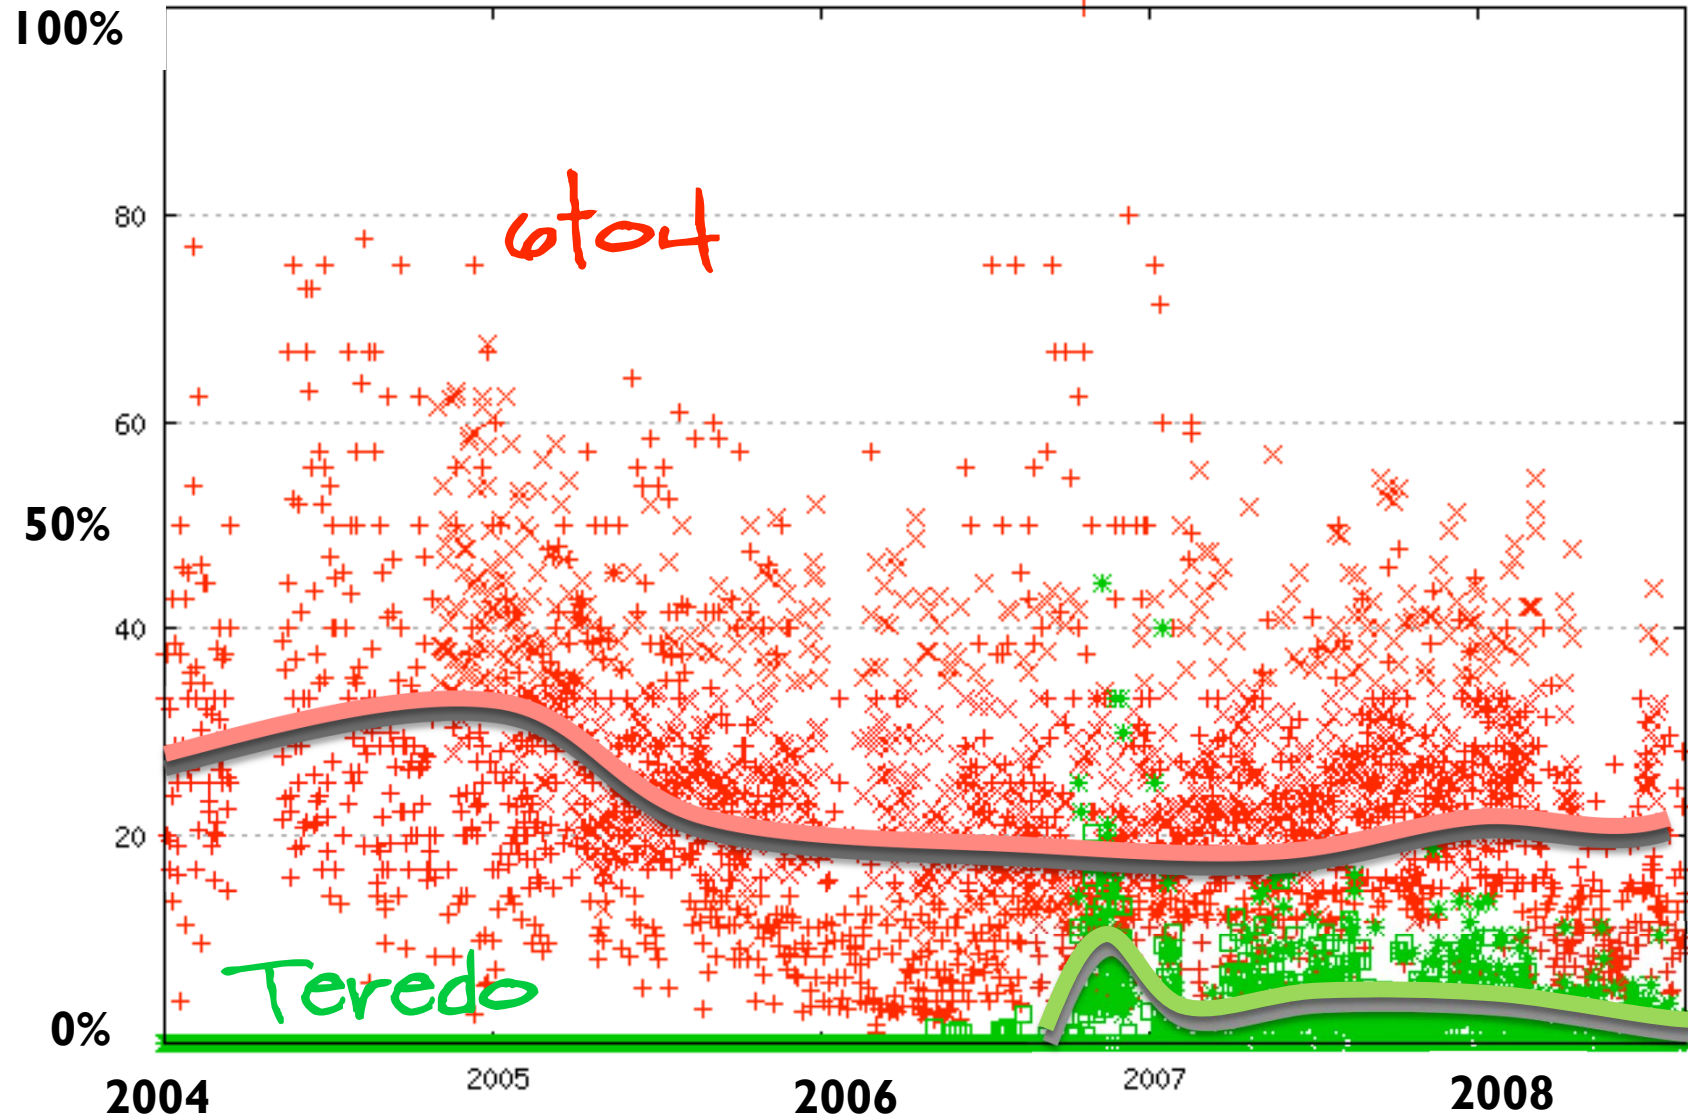
Web Server Access Statistics
Daily % of IPv6 access 1994 - today

Some Observations and Measurements

▸ IPv6 represents 0.5% of all BGP routes

▸ IPv6 is sitting at 0.5% of IPv4 in terms of host capability

Use of V6 Transition Tools

Some Observations and Measurements

▸ IPv6 represents 0.5% of all BGP routes

▸ IPv6 is sitting at 0.5% of IPv4 in terms of host capability

▸ **35% of IPv6 end host access is via host-based tunnels (6to4, teredo)**

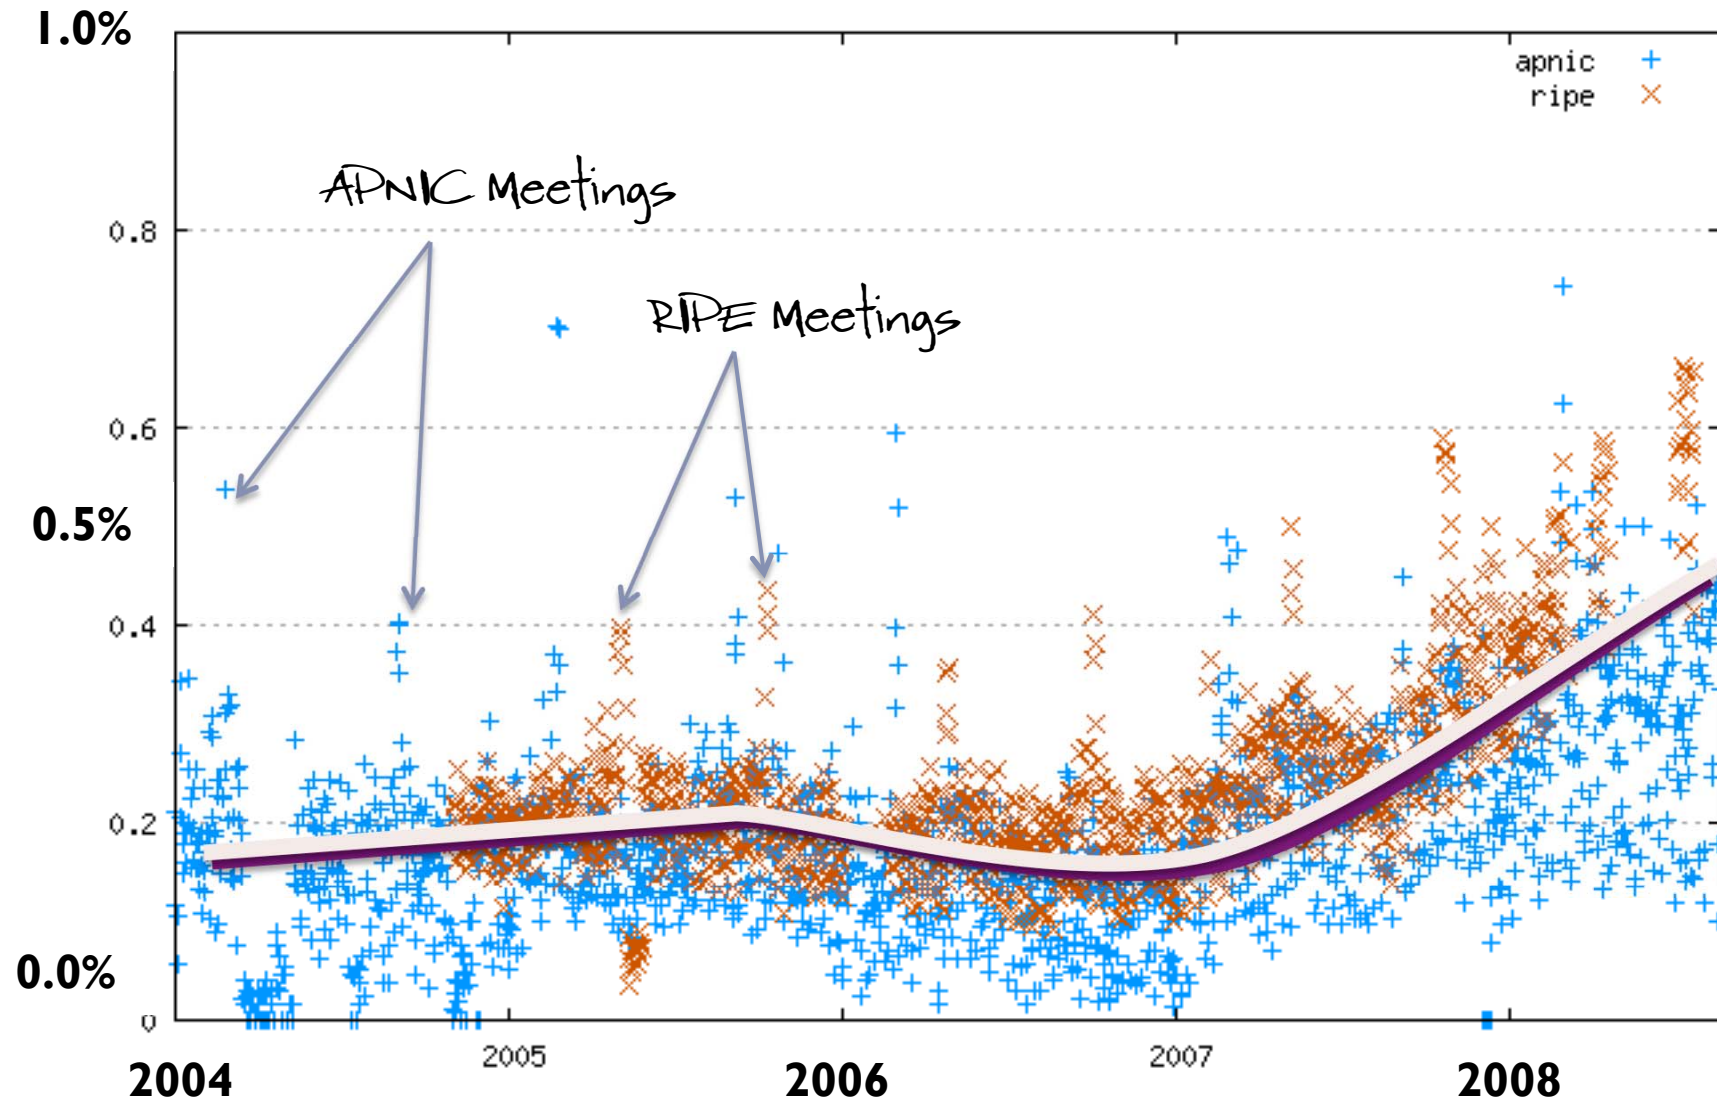# Web Server Access Statistics
Daily % of IPv6 access 1994 - today

Some Observations and Measurements

▸ IPv6 represents 0.5% of all BGP routes

▸ IPv6 is sitting at 0.5% of IPv4 in terms of host capability

▸ **35% of IPv6 end host access is via host-based tunnels (6to4, teredo)**

These last two data measurements are from a pair of relatively small web sites, strongly oriented to an IPv6-interested user base

The general number may be far smaller, and the general tunneling ratio may be far higher than that gathered from the web sites used for these measurements

AS Count IPv6 : IPv4

Some Observations and Measurements

▸ IPv6 represents 0.5% of all BGP routes

▸ IPv6 is sitting at 0.5% of IPv4 in terms of host capability

▸ 35% of IPv6 end host access is via host-based tunnels (6to4, teredo)

▸ **4% of ASes advertise IPv6 prefixes**

Some Observations and Measurements

▸ IPv6 represents 0.5% of all BGP routes

▸ IPv6 is sitting at 0.5% of IPv4 in terms of host capability

▸ 35% of IPv6 end host access is via host-based tunnels (6to4, teredo)

▸ **4% of ASs advertise IPv6 prefixes**

Actually that's a little bit misleading — here's a better summary:

15% of the IPv4 transit ASs (ISPs) announce IPv6

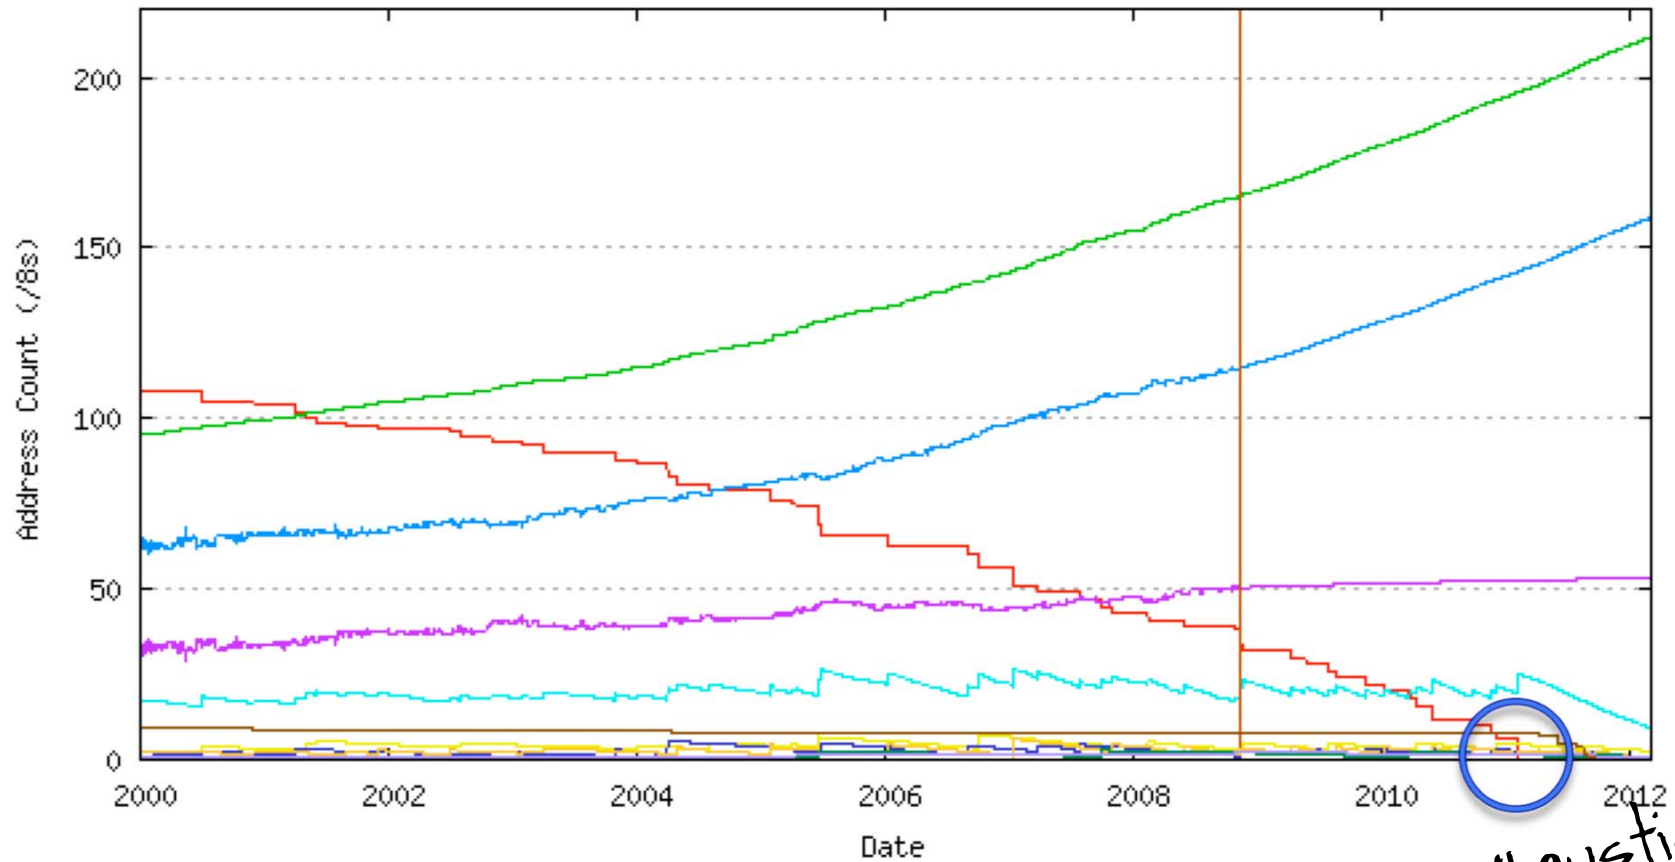2% of the IPv4 stub Ass announce IPv6

IPv4 Address Exhaustion Model

Some Observations and Measurements

- IPv6 represents 0.5% of all BGP routes
- IPv6 is sitting at 0.5% of IPv4 in terms of host capability
- 35% of IPv6 end host access is via host-based tunnels (6to4, teredo)
- 4% of ASs advertise IPv6 prefixes
- The onset of IPv4 exhaustion may occur in late 2010 – early 2011

Distribution of IPv4 address allocations 2007 - Present

Of the 12,649 individual IPv4 address allocations since January 2007, only 126 individual allocations account for 50% of the address space.
55 of these larger allocations were performed by APNIC, and 28 of these were allocated into China.
41 were performed by ARIN and 39 of these were allocated into the US

Cumulative % of allocated IPv4 address space

Cumulative % of RIR Allocations
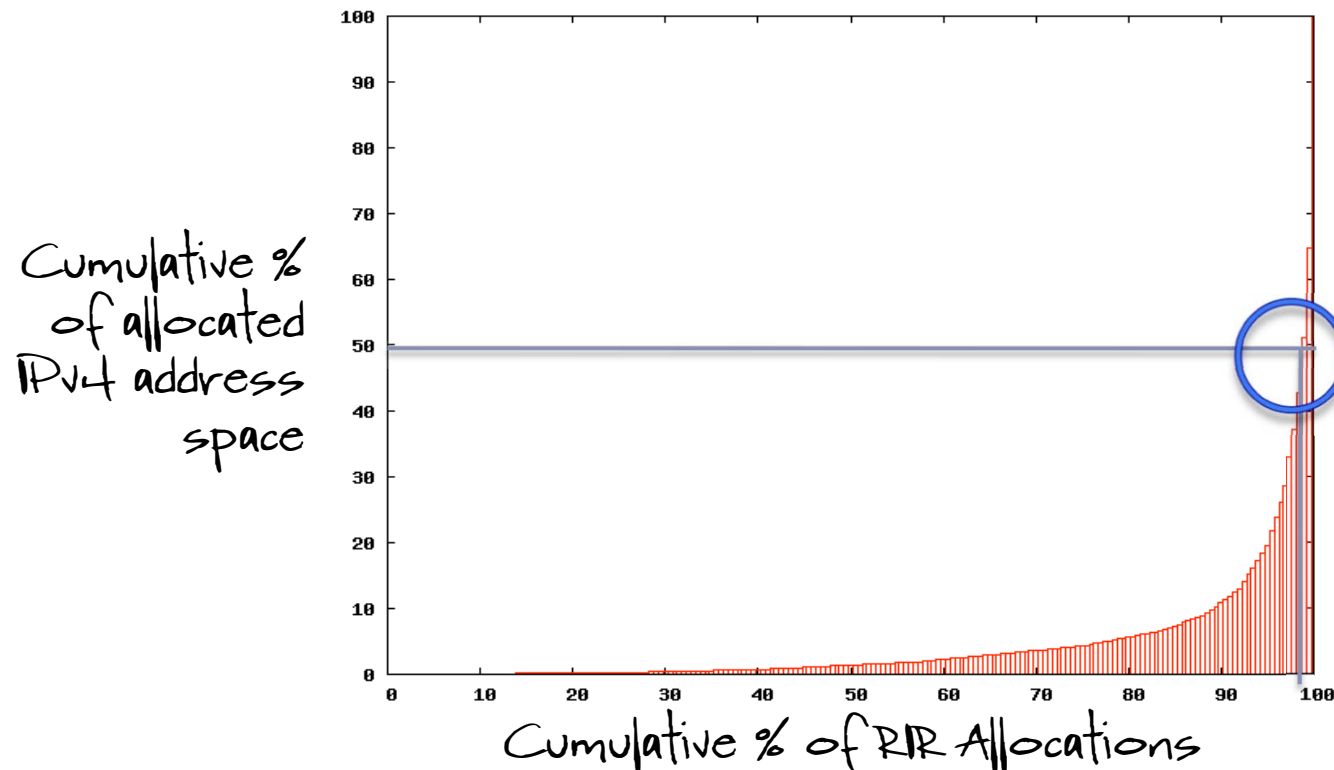
Some Observations and Measurements

▸ IPv6 represents 0.5% of all BGP routes

▸ IPv6 is sitting at 0.5% of IPv4 in terms of host capability

▸ 35% of IPv6 end host access is via host-based tunnels (6to4, teredo)

▸ 4% of ASs advertise IPv6 prefixes

▸ The onset of IPv4 exhaustion may occur in late 2010 – early 2011

▸ Large-scale capital-intensive deployments in both the developed and developing world are driving IPv4 demand today

Some Observations and Measurements

▸ IPv6 represents 0.5% of all BGP routes

▸ IPv6 is sitting at 0.5% of IPv4 in terms of host capability

▸ 35% of IPv6 end host access is via host-based tunnels (6to4, teredo)

▸ 4% of ASs advertise IPv6 prefixes

▸ The onset of IPv4 exhaustion may occur in late 2010 – early 2011

▸ Large-scale capital-intensive deployments are driving IPv4 demand

▸ We cannot avoid the situation of IPv4 demand outliving the remaining pool of unallocated IPv4 addresses

The Future Situation
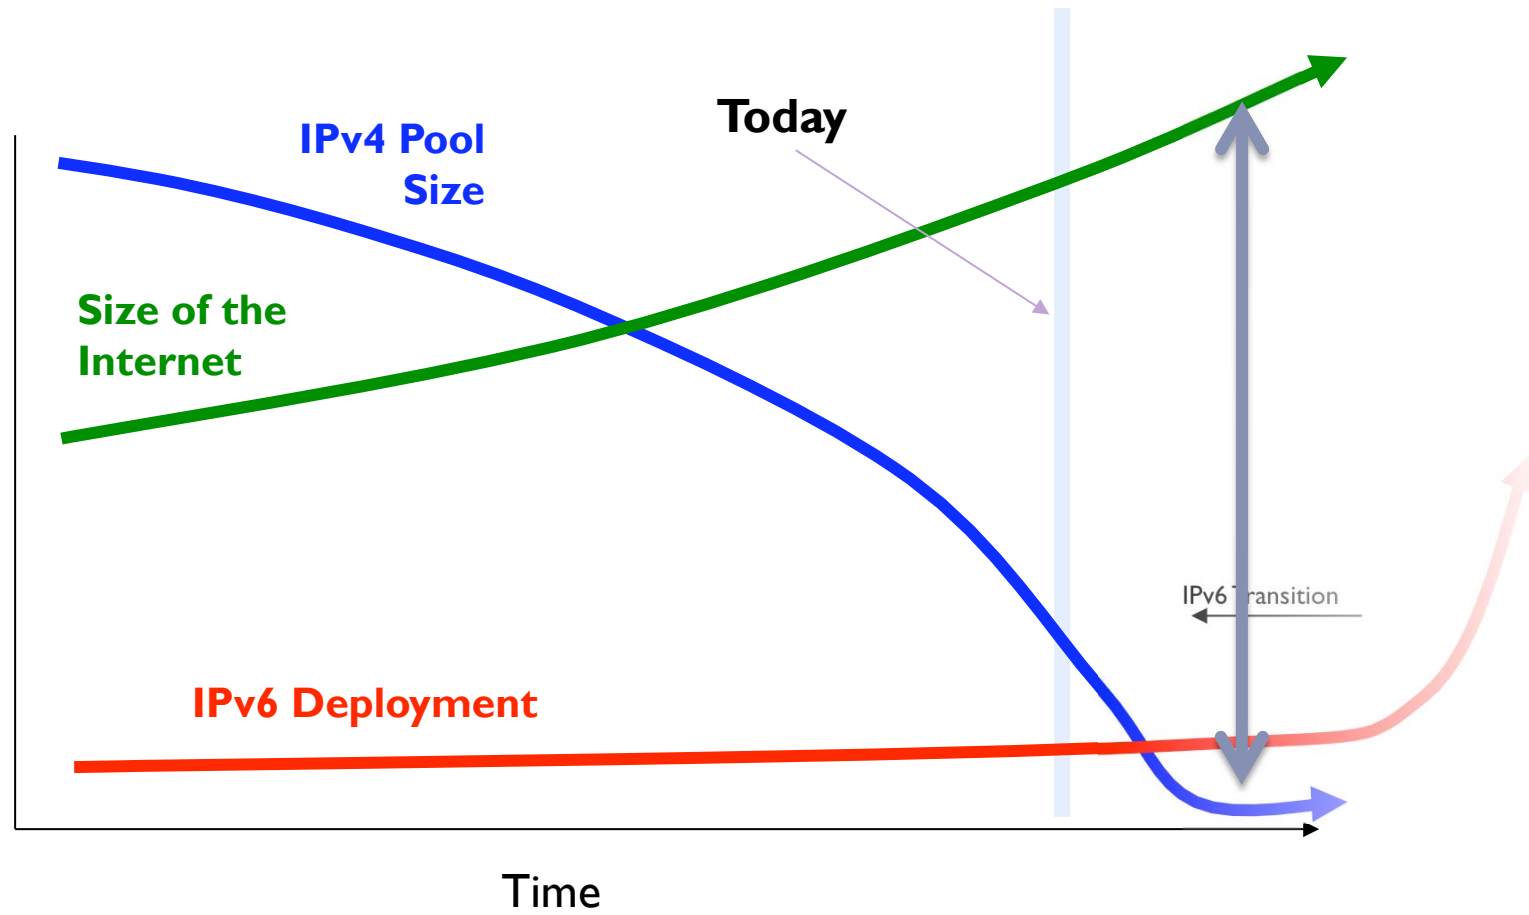
Its just not looking very good is it?

Constraints

▸ It's clear that we are going to have to use Dual Stack IPv4/IPv6 transition for some time well beyond the exhaustion of the IPv4 unallocated free pool

Constraints

▸ It's clear that we are going to have to use Dual Stack IPv4/ IPv6 transition for some time well beyond the exhaustion of the IPv4 unallocated free pool

▸ We cannot expect any new technology to assist us here in the short or medium term

Constraints

▸ It's clear that we are going to have to use Dual Stack IPv4/IPv6 transition for some time well beyond the exhaustion of the IPv4 unallocated free pool

▸ We cannot expect any new technology to assist us here in the short or medium term

▸ We are going to have to use IPv4 to span an Internet that will be very much larger than today during the final stages of this transition to IPv6

Constraints

▸ It's clear that we are going to have to use Dual Stack IPv4/IPv6 transition for some time well beyond the exhaustion of the IPv4 unallocated free pool

▸ We cannot expect any new technology to assist us here in the short or medium term

▸ We are going to have to use IPv4 to span an Internet that will be very much larger than today during the final stages of this transition to IPv6

▸ We must support uncoordinated piecemeal deployment of transitional tools, intense use of NATs and various hybrid IPv4 and IPv6 elements in the Internet for many years to come

Constraints

▸ Its also clear that the focus of any transitional effort to IPv6 will fall on the large scale deployments, and not on the more innovative small scale networked environments

Constraints

▸ Its also clear that the brunt of any transitional effort will fall on the large scale deployments, and not on the more innovative small scale networked environments

▸ We have to recognize that IPv6 is an option, not an inevitable necessity, and it is competing with other technologies and business models for a future

Challenges

▸ This is a challenging combination of circumstances:

▸ It requires additional **large-scale capital investment** in switching infrastructure and service delivery mechanisms

Challenges

- This is a challenging combination of circumstances:
    - It requires additional large-scale capital investment in switching infrastructure and service delivery mechanisms
    - There is no corresponding incremental revenue stream to generate an incremental return on the invested capital

Challenges

▸ **This is a challenging combination of circumstances:**

  ▸ It requires additional **large-scale capital investment** in switching infrastructure and service delivery mechanisms

  ▸ There is **no** corresponding **incremental revenue stream** to generate an incremental return on the invested capital

  ▸ **Displaced costs and benefits** - the major benefits of the IPv6 investment appear to be realized by new market entrants at the services and application layer rather than existing large scale infrastructure incumbents, yet the major costs of transition will be borne by the incumbent operators in the market

The Current Situation

▸ **No clear consumer signals**

  ▸ User needs are expressed in terms of services, not protocols

  ▸ No value is being placed on IPv6 by the end consumer

The Current Situation

▸ **Lack of business imperatives**

  ▸ No immediate underlying business motivation to proceed with this transition for established service enterprises with a strong customer base

  ▸ Perception that the costs and benefits of investment in IPv6 transition are disconnected

The Current Situation

- ## No clear public policy stance
  - **Uncertainty:** Having deregulated the previous structure of monopoly incumbents and encouraged private investment in communications services there is now no clear stance from a regulatory perspective as to what actions to take
  - **Risks of Action:** No desire to impose additional mandatory costs on incumbent operators, or to arbitrarily impose technology choices upon the local industry base
  - **Risks of Inaction:** No desire to burden the local user base with inefficient suppliers and outmoded technologies as a result of protracted industry inaction

What to Do?

▶ A Conservative View:

Do Nothing!

   ▶ Risk inaction for a while longer until clearer signals emerge as to the most appropriate investment direction

   ▶ Wait for early adopters to strike a viable market model to prompt larger providers enter the mass consumer market with value and capital

What to Do?

▶ A more Radical View:

Act Now!

   ▶ Take high risk decisions early and attempt to set the market
     direction with Ipv6 through leadership

   ▶ Deploy service quickly and attempt to gain an unassailable
     market lead by assuming the role of incumbent by redefining
     the market to match the delivered service

Further Thoughts

▶ A Public Sector Regulatory View

Think about it some more!

  ▶ Its about balance, efficiency and productive private and public sector infrastructure investments that enable leverage to economic well-being

  ▶ Its about balance between:

    ▶ industry regulatory policies for the deployment of services to meet immediate needs of local users and local industry, with

    ▶ public fiscal policies to support capital investments to sustain competitive interests in the short term future, with

    ▶ economic developmental policies to undertake structural investments for long term technology evolution

What to do?

▸ **What can we do about this transition to IPv6?**

  ▸ Is the problem a lack of information about IPv4 and Ipv6? Do we need more slidepacks and conferences to inform stakeholders?

  ▸ Should we try to energise local communities to get moving?

  ▸ Should we try to involve the public sector and create initial demand for IPv6 through public sector purchases?

  ▸ Should we try to invoke regulatory involvement?

  ▸ Should we set aspirational goals?

  ▸ Should we attempt to get the equipment vendors and suppliers motivated to supply IPv6 capability in their products?

  ▸ Should we try to invent new transitional technologies?

  ▸ Or should we leave all this to market forces to work through?

What to do?

- Maybe this is not an accidental problem
- Maybe the shortcoming lies in the architecture of IP itself
- And maybe this situation represents an opportunity to do something about it

I have a couple of my own modest suggestions about what to do, as a result of these considerations ...

Today's Agenda

1. Get moving on today's issues

Operational Tactics:
Tomorrow's Dual Stack Internet

▸ Can we leverage investments in IPv6 transitional infrastructure as a 'natural' business outcome for today's Internet?

▸ How do we mitigate IPv4 address scarcity? By attempting to delay and hide scarcity or by exposing it as a current business cost?

▸ Do we have some viable answers for the near term? Do the emerging hybrid V4/V6 NAT models offer some real traction here in terms of scaleable network models for tomorrow's networks?

▸ What's the timeline to deployment for these hybrid NAT approaches?

More Agenda Items for Today

1. Get moving on today's issues

2. And do not forget about tomorrow

Overall Strategy

▸ How do we evolve our current inventory of wires, radios and switches into tomorrow's flexible and agile network platforms to allow for innovation in services to meet the demand of an increasingly diverse application portfolio?

▸ Or should we consider more capable applications layered across a heterogenous network substrate?

Overall Strategy:
Where is this leading?

▸ What's the research agenda?

▸ What can we learn from this process in terms of architectural evolution of networking services?

▸ What's **really** important here?

  ▸ IPv6?

  ▸ Or a service evolution that exploits a highly heterogenous networked environment?

  ▸ Why do today's services need protocol uniformity in our networks?

  ▸ Can we build a stable service platforms using hybrid IP protocol realms?

One evolutionary view of network architecture – moving up the stack

circuit networking - yesterday

shared capable network with embedded applications

simple 'dumb' peripherals

single simple application

packet networking - today

simple datagram network

complex host network stacks

simple application model

identity networking - tomorrow

realms of simple datagram networks

locator—based simple host network stacks

identity—based complex application overlays

Where Next?

▸ Perhaps all this is heading way further than just IPv6

▸ Perhaps the real opportunity here is about breaking away from the two-party communications model as an overlay above a uniform protocol substrate and looking at a model of peer-networking application architectures with relay and rendezvous agents layered on a heterogenous base

▸ Perhaps we are starting to work on the challenges involved in a new generation of identity-based networked services as a further evolutionary step in networking service architecture

Thank You