The Internet in Transition:

The State of IPv6 in Today's Internet

Geoff Huston

Asia Pacific Network Information Centre

The Internet...

has been a runaway success that has transformed not just the telecommunications sector, but entire commercial and social systems are being transformed by the Internet

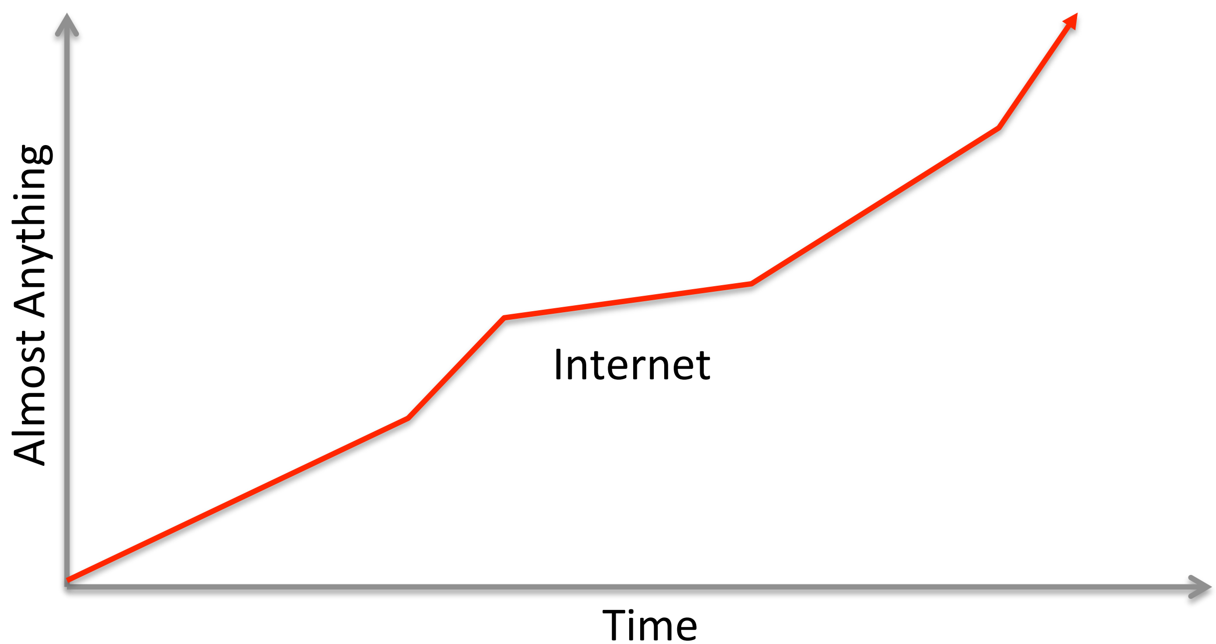

Growth Pressures

The protocol was designed in an era of mainframe computers, where the largest networks of the day connected 100’s of devices

The same protocol is being used today in a context of use that spans billions of devices

We confidently anticipate further growth

Scaling Critical Infrastructure

- Silicon-based equipment scales with Moore's Law
  - As long as the aggregate growth rate is below doubling every two years economies of scale still hold in this area
- Names scale within the structure of a loose hierarchy
  - Adding name names at the leaf points of the name structure scales at a level of O(n)
- Addresses are fixed size elements in the protocol
  - And this is a problem, and has been recognised as a problem for more than 20 years
  - It's now an urgent issue because of the exhaustion of IPv4 addresses in AsiaPac, Europe and the Middle East

Contact us | Background | APNIC Research & Development | APNIC website

Labs.APNIC.NET - IPv4 Address Allocation Report

Report Date: 03-Dec-2012 01:51 UTC.

IPv4 Unallocated Address Pool Exhaustion:

03-Feb-2011

Projected RIR Address Pool Exhaustion Dates:

RIR	Projected Exhaustion Date	Remaining Addresses in RIR Pool (/8s)
APNIC:	**19-Apr-2011** (actual)	0.8988
RIPE NCC:	**14-Sep-2012** (actual)	0.9570
ARIN:	**05-Jun-2014**	3.0990
LACNIC:	**19-Sep-2014**	2.9664
AFRINIC:	**01-Jan-2022**	3.9244

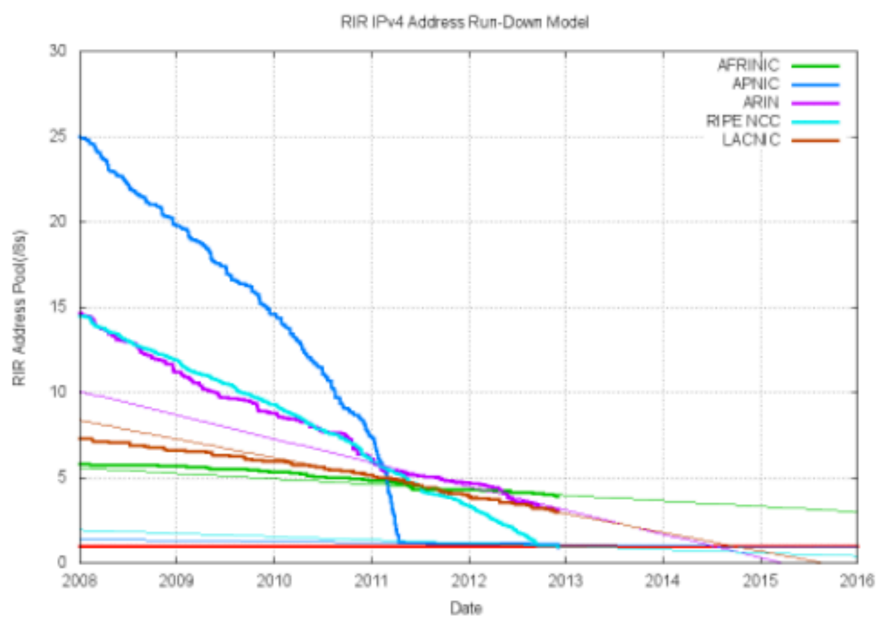

Projection of consumption of Remaining RIR Address Pools

Extending IP Addresses

IPv6

New protocol header

- lengthen the address field by adding more bits to the packet header
- Preserves the architecture of the network
- Issues about backward compatibility

NATs

"Share" an address across multiple users

- Use the transport protocol bits to share a single address
- Preserves the application functionality of the network
- Backward compatible
- Destroys the architecture of the network

IPv6

- Protocol defined in the mid-90’s
- Reference open source implementations available mid-late 90’s
- Now implemented for most device platforms and enabled for use (Microsoft Windows, Apple OSX and iOS, Android, Linux, ...)

Deploying IPv6

- All devices need to be reprogrammed to include an IPv6 stack in addition to an IPv4
- Infrastructure elements need to be re-configured to include IPv6 access as well as IPv4
- Access networks and CPE need to be re-configured/replaced to support IPv6 as well as IPv4

Deploying IPv6

- All devices need to be ... to include an IPv6 stack in addition to an IPv4

  Well underway – more than 50% of devices now support IPv6

- Infrastructure elements ... to include IPv6 access as well as IPv4

  Underway - transit services, DNS services, content services being reconfigured

- Access networks and CPE need to be re-configur... to support Ipv6 as well as IPv4

  This is a potential problem area!

IPv6 capability, as seen by Google

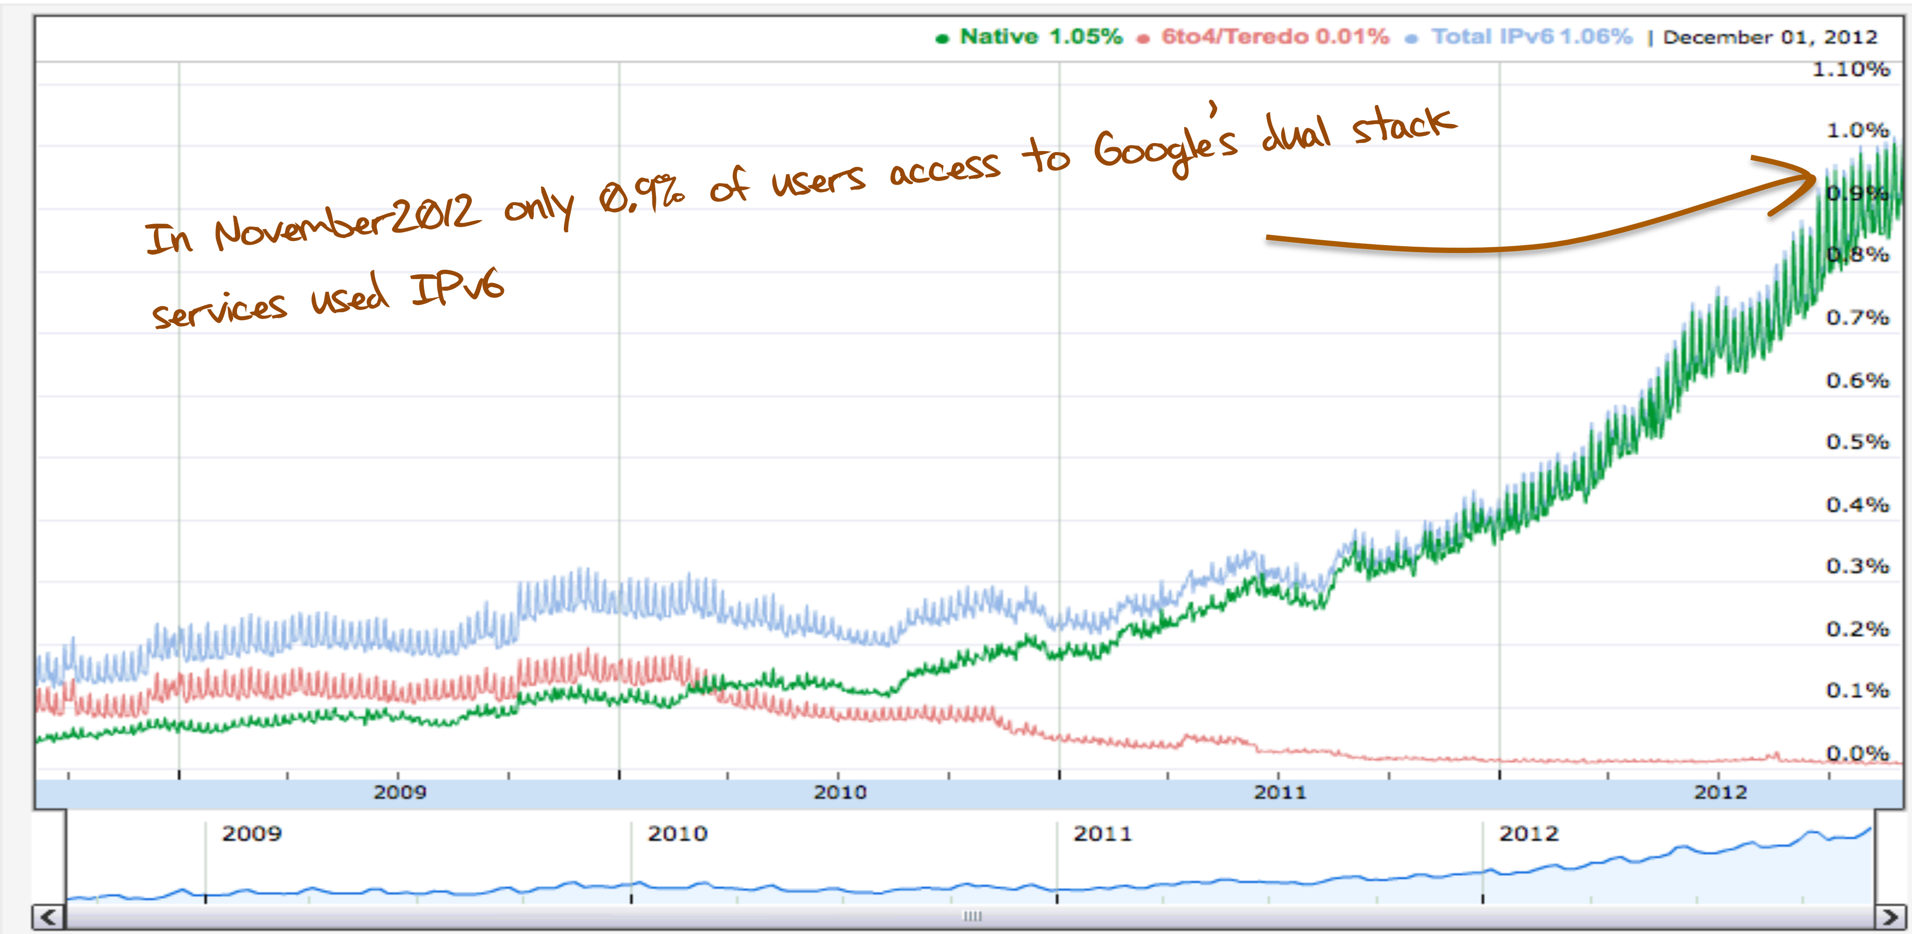

Source: http://www.google.com/intl/en/ipv6/statistics/

IPv6 capability, as seen by APNIC

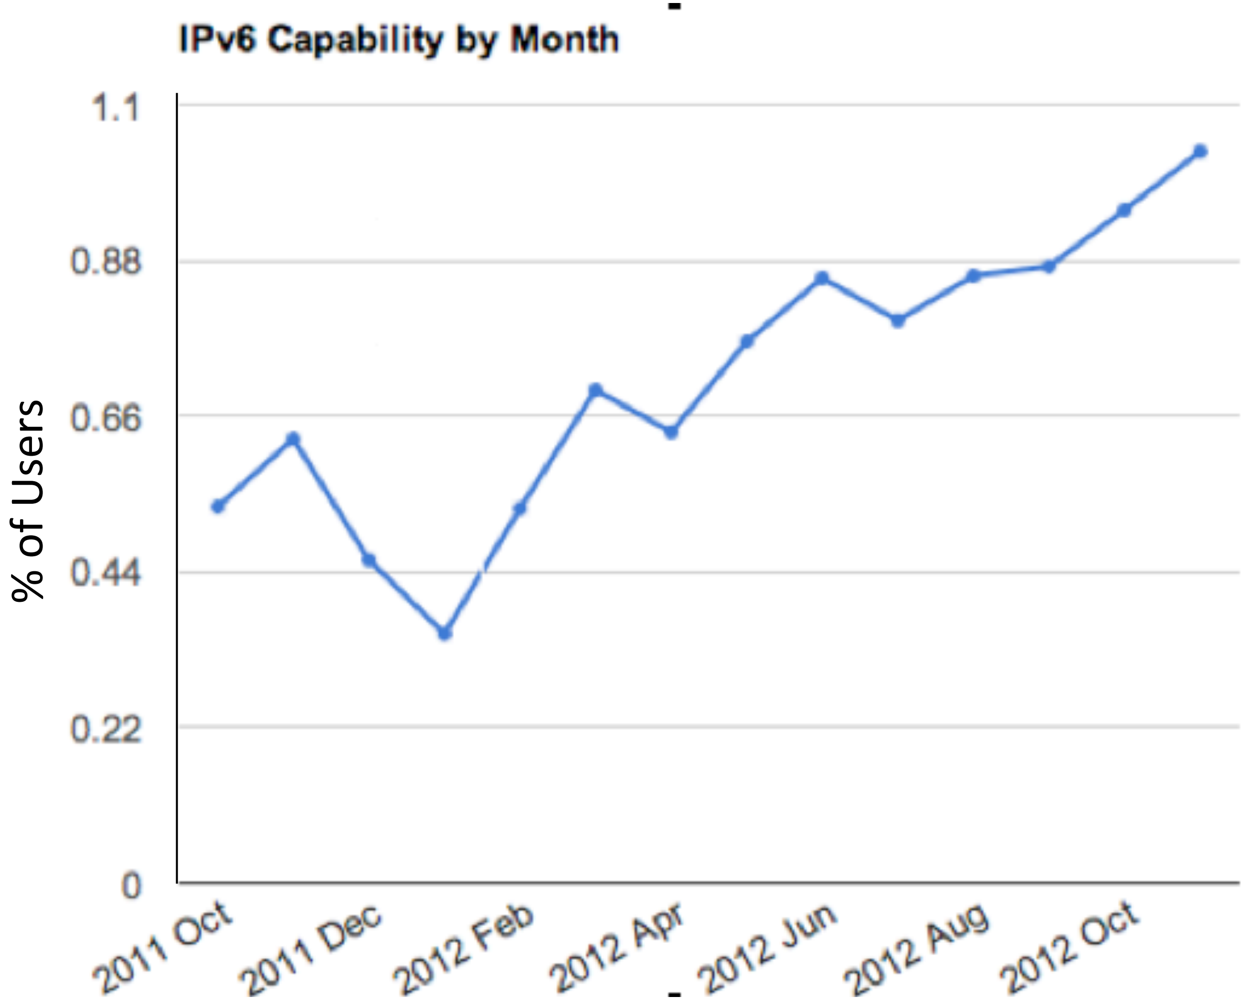

Source: http://labs.apnic.net/ipv6-measurement/Regions/001%20World/

Where is it?

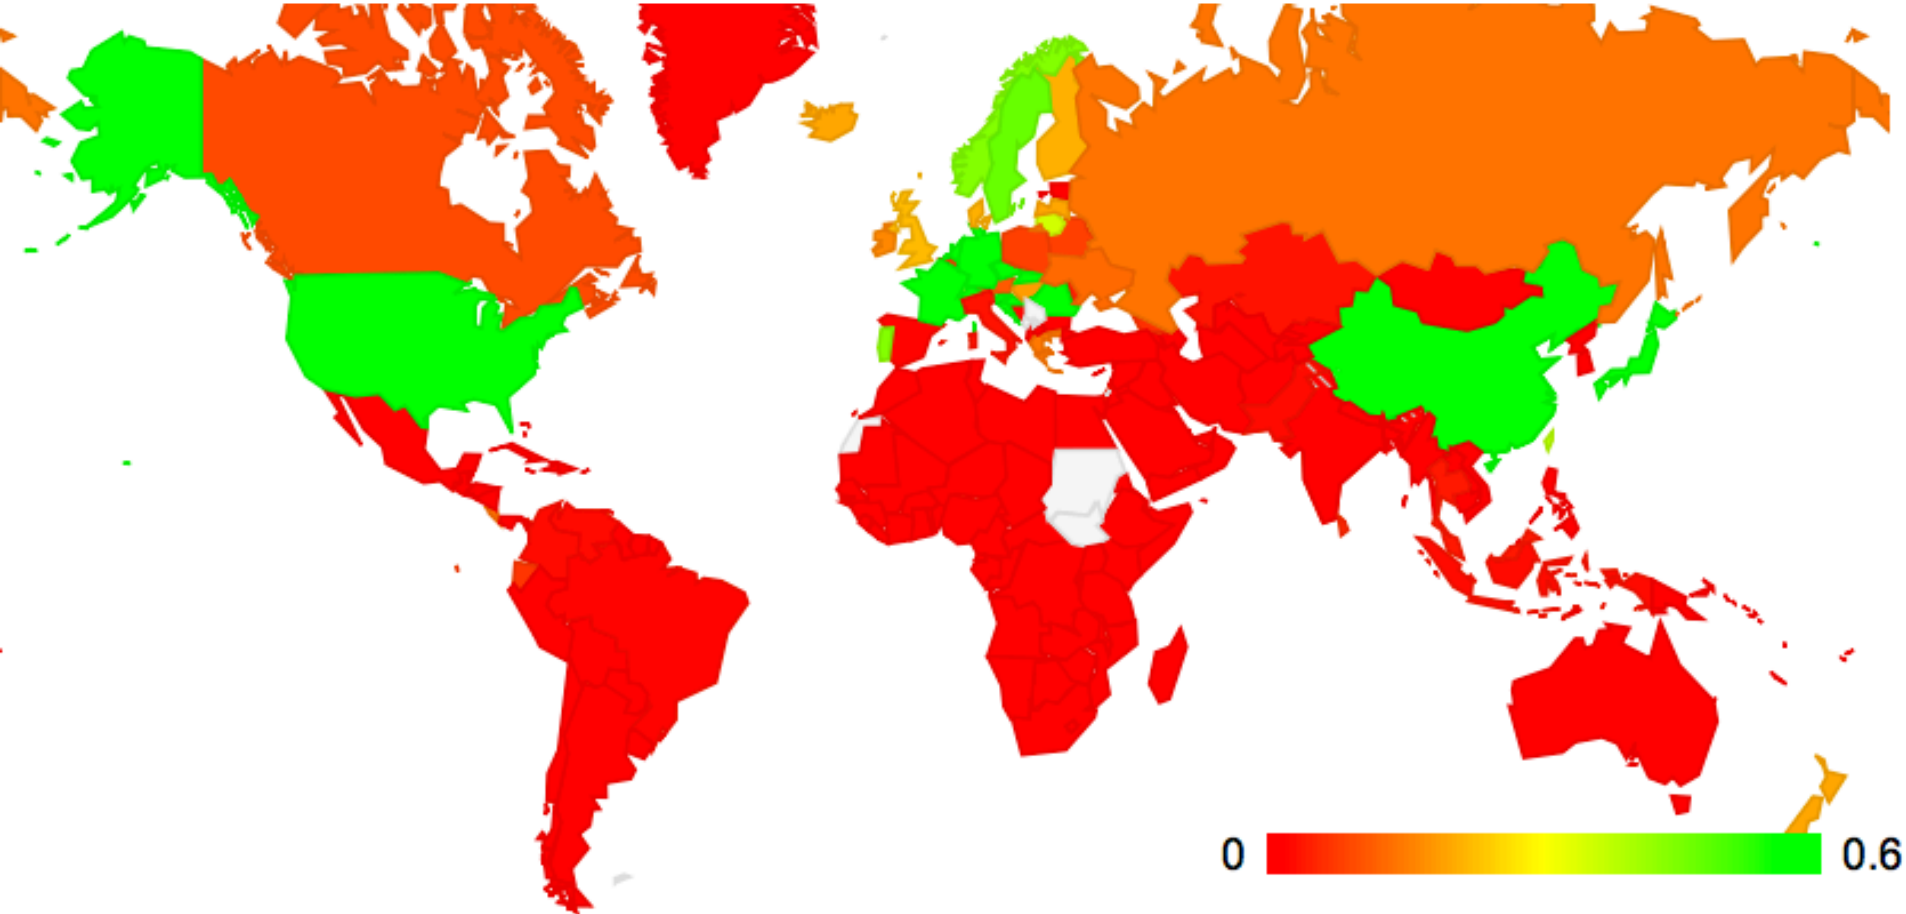

% of users preferring IPv6 – per country

http://labs.apnic.net/index.shtml

Why is IPv6 not happening?

The major issue appears to be in the business structure of the "last mile" access networks

The usual business incentives that would drive investment in new services appear to be lacking for IPv6 – IPv6 represents cost without benefit for many access providers

What happened 20 years ago?

If IPv6 is such a problem today then how did this industry adopt IPv4 in the first place?

PSTN Circuits to IP Packets: The Demand Schedule Shift

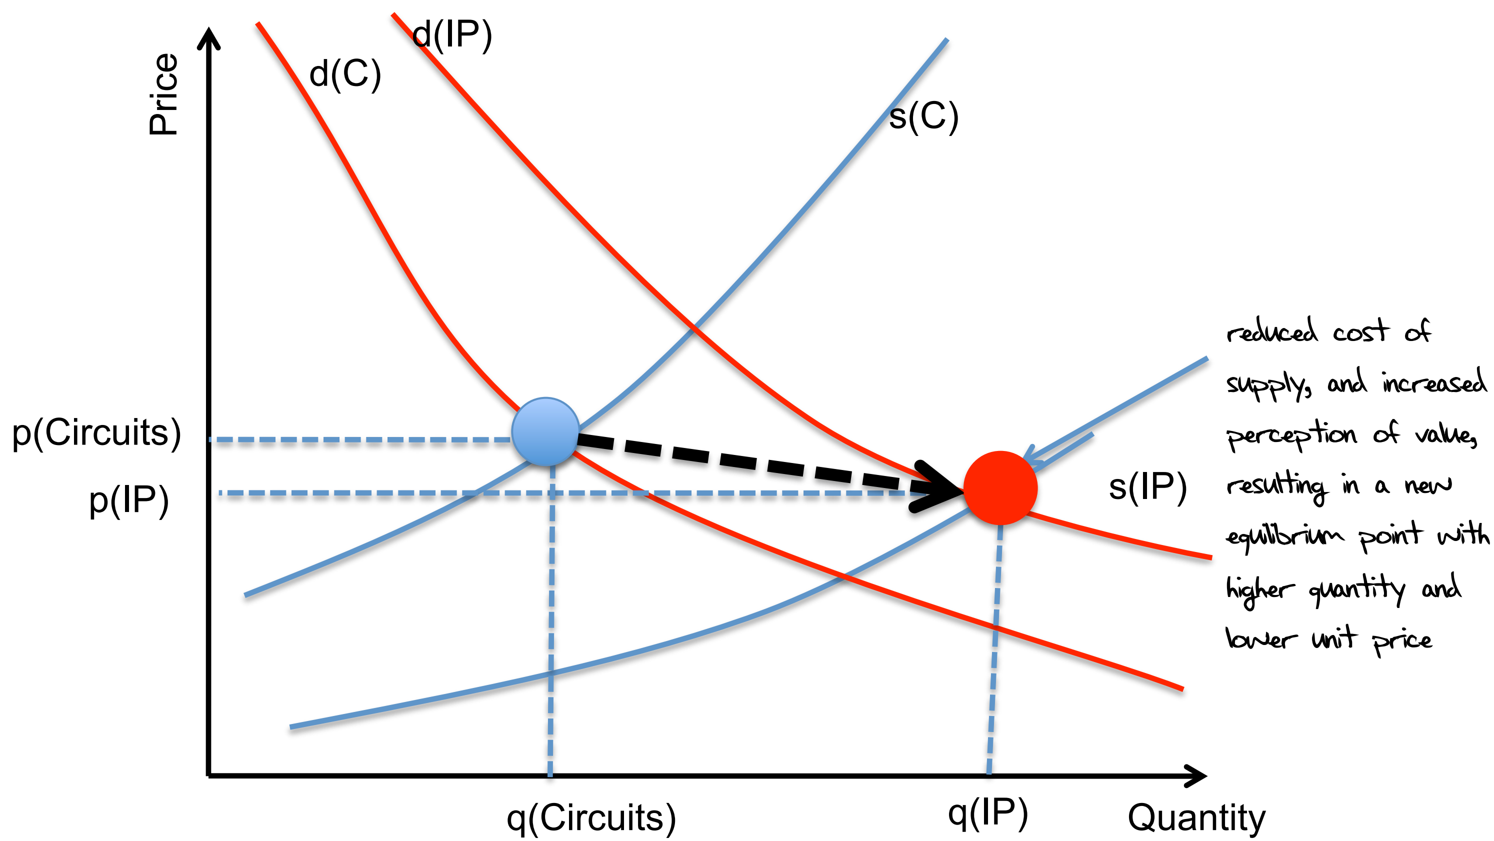

IPv6 vs IPv4

Are there *competitive differentiators*?

✗ $cost_4 = cost_6$

✗ $functionality_4 = functionality_6$

no inherent consumer-visible difference

no visible consumer demand

no visible competitive differentiators other than *future risk*

IPv4 to Dual Stack: The Demand Schedule Shift

Supply side cost increase due to Dual Stack operation

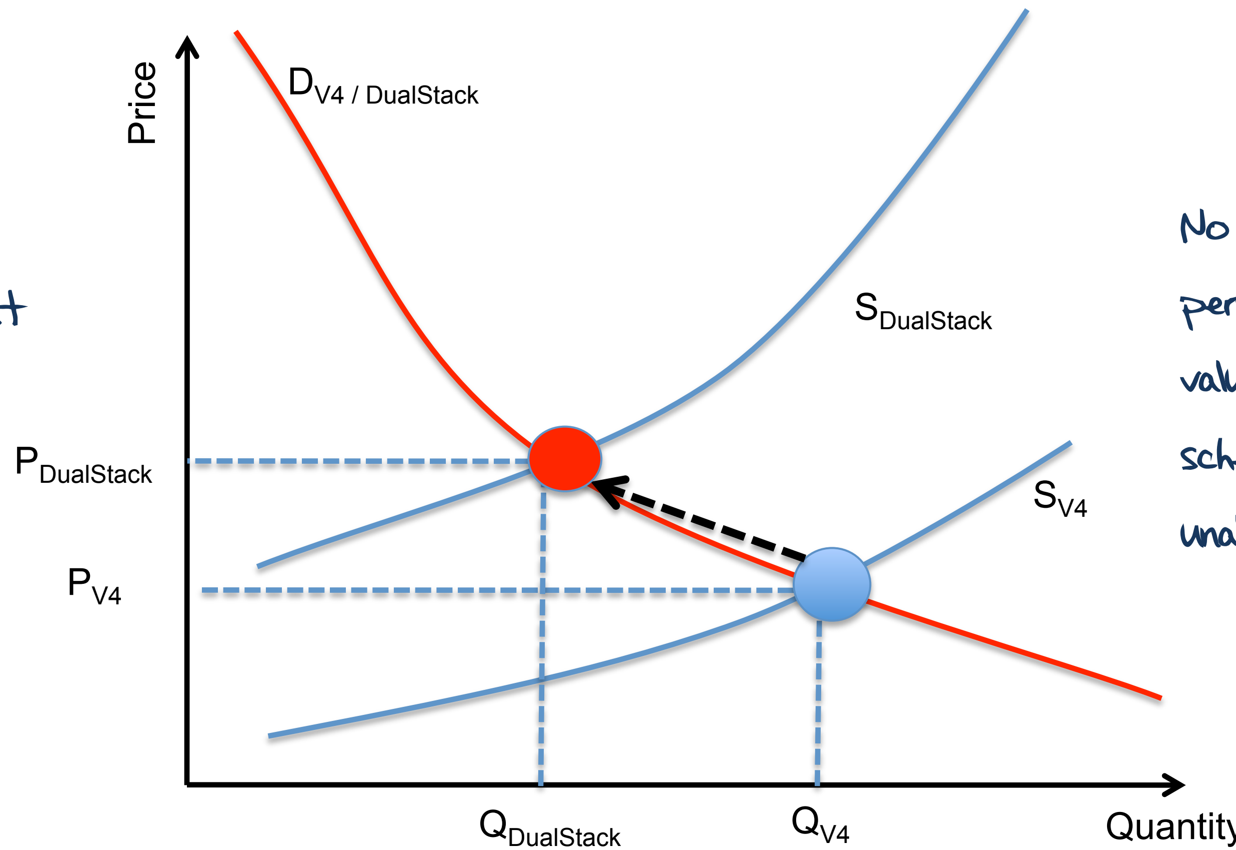

No change in perception of value, so demand schedule is unaltered

Equilibrium point is at a lower quantity if Dual Stack supply costs are passed on to customers

The Transition to IPv6

- Given that
  - we've left it so late in terms of the scale of the transition
  - the degree of difficultly with IPv4 exhaustion
  - there appears to be little economic motivation from the carriage side of the industry to embark on this transition
- Will market forces actually drive the industry to adopt IPv6 at all?

The Alternative to IP6

Increase the density of NATs by adding CGNs to the carriage infrastructure

- CGNs share a single address across multiple customers by multiplexing on the transport port addresses
- This can be achieved incrementally, with modest outlay, and without altering the customer's equipment or applications, and without coordination with any other provider or content delivery system
- With IPv4 exhaustion this is a forced decision – once a SP runs out of IPv4 addresses this is a cost effective option to support further growth

IPv4 to CGNs: The Demand Schedule Shift

Supply side cost decrease due to CGN operation offset by opportunities for leverage over content

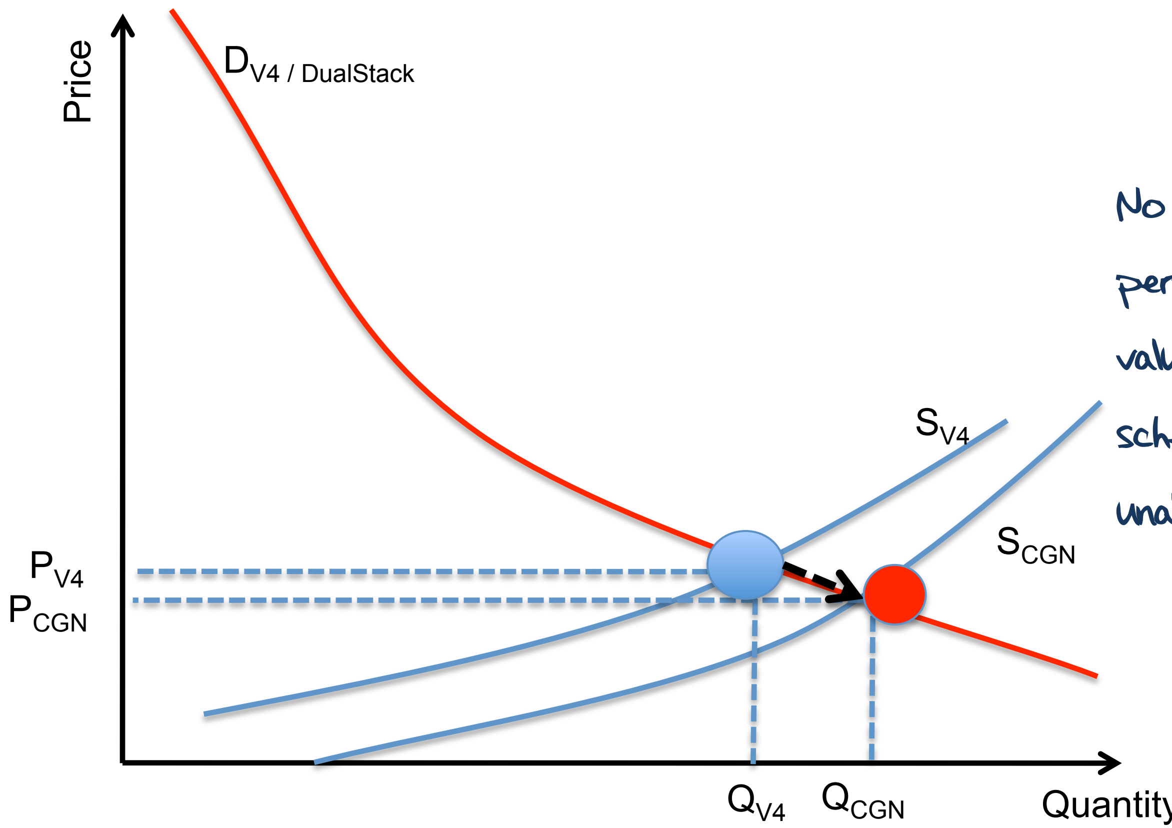

No change in perception of value, so demand schedule is unaltered

Based on leverage over content CGNs may produce a preferred outcome for the access provider market

Carriage vs Content

The architecture of the Internet struck a new balance between carriage and content:

- Content no longer required the permission of the carriage providers
- Any form of content, delivered in any fashion that optimized the efficiency of the user's interaction with the content could be implemented on the Internet
- The carriage network was unaware of the nature of the content and service transactions

The value of the "Internet Economy" is the value of this redefinition of the provision of goods and services, and the removal of carriage level impositions and overheads from the picture

Carriage vs Content

Carrier NATs in IPv4 fundamentally change this balance:

- Carriage providers have direct visibility of all user transactions
- Carriers can directly alter the quality of the service delivered to users for individual services through manipulation of CGN behavior
- Carriers can directly create barriers of access to users, forcing content providers to pay an access premium for direct access to the carrier's user base
- There is no efficient alternative for content to access users given the address exhaustion issue and the unique local monopoly position of access providers

CGN Risks

The CGN approach was intended to be a stopgap measure for IPv4 address exhaustion

But there are long term risks here:

- The major risk is that the incumbent content providers join with the incumbent carriers to exploit this situation to create:
  - elevated barriers to entry for new content
  - limitations on the forms of innovation for content delivery
- Incumbents in carriage and content are then in a unique position to define the terms and conditions for future competition
  - This may result in a small number of actors with overarching control of carriage and content over the entire communications system

Why IPv6?

IPv6 represents the most efficient path to support an open network that can sustain efficient competitive access to the carriage and content service roles

And efficient competitive access to all parts of this activity underpin almost all of the expectations of future efficient growth of the Internet Economy

The 2008 Seoul Declaration for the Future of the Internet Economy

"Encourage the adoption of the new version of the Internet protocol (IPv6), in particular through its timely adoption by governments as well as large private sector users of IPv4 addresses, in view of the ongoing IPv4 depletion."