

Hello!

Förlåt

Sorry

Hej!

Hello

Hur många

How many

IPv6

IPv6

presentationer

presentations

har du sett

have you sat through

de senaste 10 åren?

in the last 10 years?

20?

do you really want me to translate this?

200?

c'mon - I'm not doing this one

2,000?

or this

Har du inte sett nog?

Had enough yet?

Vill du verkligen

Do you really

se

want to sit through

en till

yet another

tröttsam

mind numbing

presentation

presentation - really!

om

about

hur IPv6 kommer bli

how IPv6 is going to be

större

bigger

bättre

`better`

snabbare

`faster`

och vackrare?

and shinier?

Inte jag heller.

Neither do I

Låt oss prova

So lets try

någonting annat.

something else.

ok?

ok?

Efter 10 års

After 10 years

väntan

of waiting

på ett IPv6

for an IPv6

Internet

Internet

har vi inte lyckats med

`we've achieved`

någonting.

nothing

Så

So today

låt oss börja

let start off

med ett annat ord

with another word

Misslyckande!

I really hope that's the swedish for: failure!


Usual disclaimer stuff:

All the bad ideas here are mine (Geoff).

Any good ideas probably came from someone else (Patrik)!


Today

IPv4 Address Pool Status

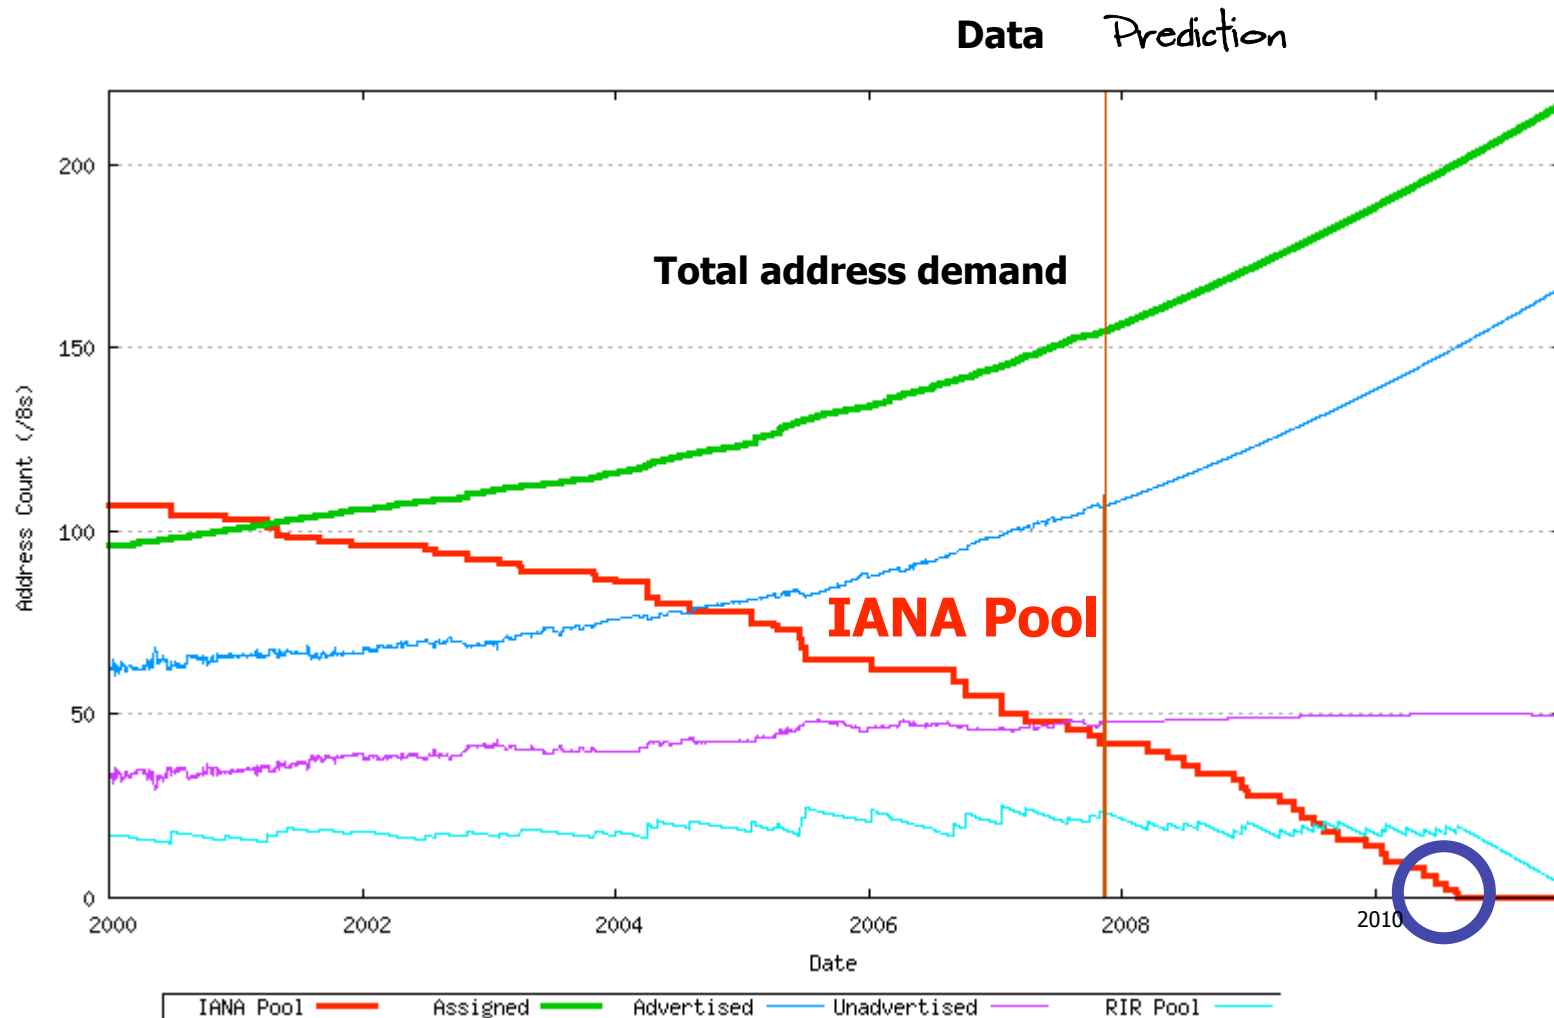


Tomorrow

Tomorrow



Ooops!




That's 15th November 2010

<http://ipv4.potaroo.net>

That's 15th November 2011

YAWN - seen all this before
Its now **BORING!**



<http://ipv4.potaroo.net>

That's a highly uncertain prediction
– it could be out by as much as 18
months



We can't model changes in demand due to:

Panic - last minute rush

New Policies – “reservations” of remaining address space

Change of relative Ipv4 / IPv6 demands

And modeling uncertainty due to:

highly skewed data used to make projections



Let's say some time between
late 2009 and early 2011




what then?




IPV6!

We had this plan ...





What's the revised plan?




If IPv6 is the answer then...



Plan A: its time to move!

The global internet adopts IPv6 universally before **January 2009** and completely quits all use of IPv4 well before address pool exhaustion occurs



If IPv6 is the answer then..




Plan A: **its time to move!**

The global Internet, with more than 1.7 billion users, a similar population of end hosts, and hundreds of millions of routers, firewalls, and billions of lines of configuration codes, and hundreds of millions of ancillary support systems, where only a very small proportion are IPv6 aware, are all upgraded and fielded to work with IPv6 **in the next 60 days**, and then completely quits all use of IPv4 in **10 days later**.



Really!



BIG and FAST don't go together!

If IPv6 is the answer then...



Plan B: **Dual Stack**

Leisurely IPv6 deployment

and

Persist with IPv4 networks using more NATs

If IPv6 is the answer then...



Plan B: **Dual Stack**

Make IPv4 work using more intense levels of NAT deployment in new products and services for as long as the existing deployed networks continue to use Ipv4

This may take a decade or two!

If IPv6 is the answer then...



Plan B: **Dual Stack**

So if IPv4 is a necessity for the next 10 or 20 years, what exactly is IPv6's role here?

What immediate marginal benefit is obtained from the additional cost of deploying IPv6 in a dual stack?



Its just not looking very good is it?

Why are we here?

Its Just Business ...



This entire network is customer funded

Every vendor is intensely focussed on meeting customer needs

Customers have absolutely no clue what this IPv6 stuff is about - so they are not paying extra for IPv6!

And vendors and service providers are not about to build IPv6 for free

We appear to be seriously wedged!

Or Business Failure?



IPv6 adoption offers all the marginal benefit of a pretty minor technology change with all the costs and disruption of a major forklift upgrade

On the other hand there are more options...

What options for the Internet's future exist that do not necessarily include the universal adoption of IPv6?

The Failure Option



What if IPv6 doesn't happen?

The Failure Option



What if IPv6 doesn't happen?

Existing network deployments continue to use IPv4
- no change there

New networks will have to use IPv4 – no change
there either

We are going to have to make IPv4 last past
exhaustion, coupled with intense use of NATs –
no change there either!



If IPv6 is NOT the answer then...

Plan C: **IPv4 for ever**

~~Leisurely IPv6 deployment
and~~

Persist with IPv4 networks using more NATs

Making IPv4 Last Longer



Redeploy “idle” IPv4 addresses?

Not every address is “in use”

End host utilization levels of addresses are estimated to be around 5% - 20% of the address pool

So could we flush more addresses back into circulation?

Yes, but it will take money and markets to flush them out!

NATs on Steroids?



We need to get **really** good at NATs ...

Fun new products to play with: carrier grade NATs?

Multi-level NAT deployments both at the customer edge and within the ISP network

Standardise NAT behaviours to full cone behaviour allow application determinism and maximum address / port utilization

Load applications with greater levels of context discovery, multi-party rendezvous, and adaptive parallelism

NAT Futures



Are NATs just more of the same? Is this the “safe” option?

How far can NATs scale?

How complex can we get with this network?

Are we willing to find out?

NAT limits?



Recent studies on application behaviour:

Applications use parallel sessions to improve performance

Each host needs an allowance of 100 – 300 ports for the more extravagant applications

Each NAT IP address can serve 200 hosts, or maybe 100 customers within the framework of existing application behaviours - without creating too much havoc!

Numbers, numbers, numbers



Assume that:

dual stack transition will take a further 10 years

the growth pressure for network connectivity will average 200 million new connections per year

All growth will be using IPv4

A /16 could service around 6 million customers if you achieved 100% packing density with NATs

Numbers, numbers, numbers



Assume that:

dual stack transition will take a further 10 years

the growth pressure for network connectivity will average 200 million new connections per year

All growth will be using IPv4

CGNats achieve average of 50% address utilization efficiency with allowance of 600 ports per customer

Could that scale to 1 billion customers on a /8 ?

Numbers, numbers, numbers



Assume that:

dual stack transition will take a further 10 years

the growth pressure for network connectivity will average 200 million new connections per year

All growth will be using IPv4

CGNats achieve average of 50% address utilization efficiency with allowance of 600 ports per customer

Then the IPv4 requirements for the next 10 years of Internet growth would be possible within a pool of 4 /8s !

But what about the next 10 years?

And the next 10?

And ...



Maybe that's pushing NATs a bit too far

what other options do we have?



If IPv6 is NOT the answer then...

Plan X: **end-to-end IP is NOT the answer
either!**

huh?

Application Level Gateways!



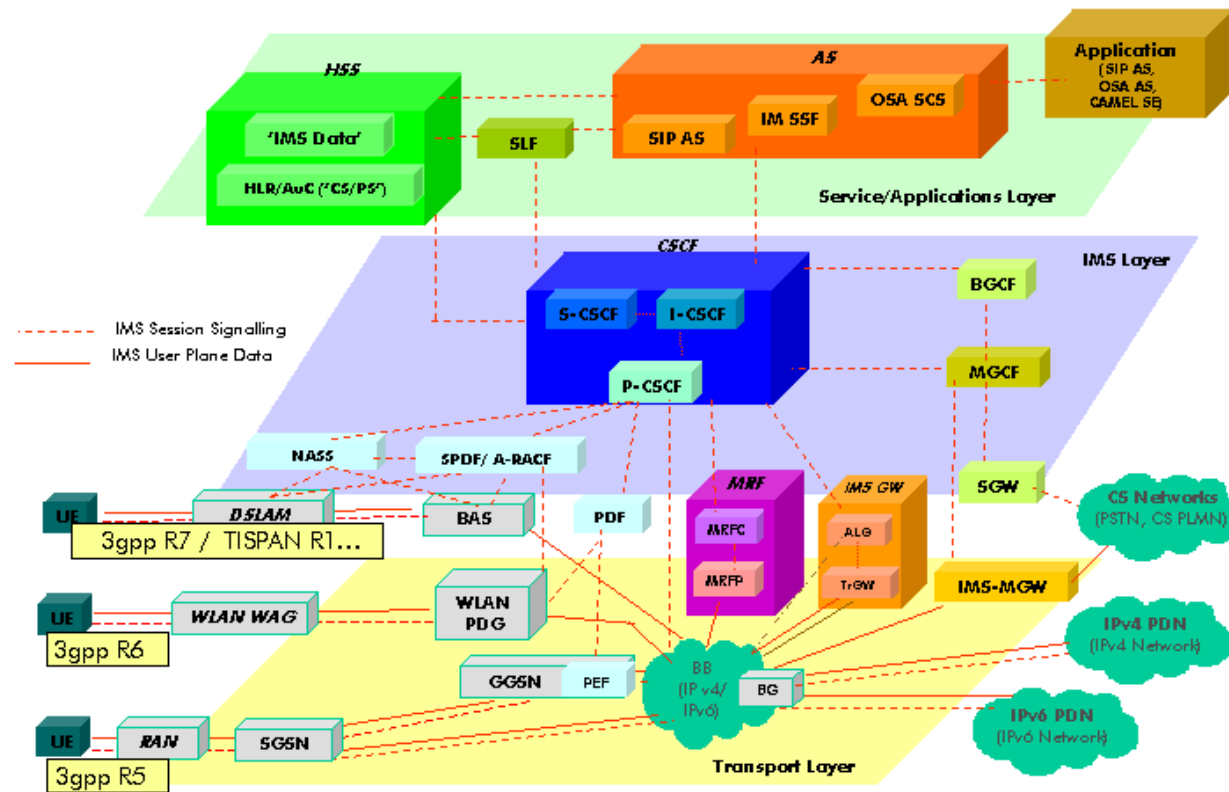
For example:

Use the 3G approach - IMS

IMS is an architecture of application level gateways

- front-end proxies act as agents for local clients
- applications are relayed through the proxy
- no end-to-end IP at the packet level

Yes, it's VERY ugly!



And IPX also smells...

Maybe it gets even uglier!

“The true technical solution to the challenge of convergence comes as we make the move to IMS, or IP Multimedia Subsystems, which will provide the common control and protocols for applications to work across our networks. We’ve been involved in the push for IMS since its inception. In 2006, we drove an initiative called “Advances in IMS”, which was executed by a task force of companies, whose purpose was to catalyze closure on worldwide standards for IMS which would make its deployment pragmatic in the near-term for operators. I’m happy to say that we succeeded. With IMS, the customer will no longer be stranded on separate islands of technology for things like messaging, voice, or video. Instead, we’ll be able to build an application once and have the network deliver it to customers wherever they need it.”

Dick Lynch CTO Verizon, 20 August 2008

But is something deeper about networking architecture evolution lurking here behind the ugliness?



circuit networking

shared capable network with embedded applications
simple 'dumb' peripherals

packet networking

simple datagram network
complex host network stacks
simple application model

identity networking?

sets of simple datagram networks
locator-based host network stacks
identity-based application overlays



Are you feeling lucky today?

Do you understand enough about this to bet the entire future of the Internet on this theory of the evolution of network architectures?

Probably Not!

So – I hope we have shown you that IPv6 is not necessary outcome, and that there are failure options.

And some sectors of this industry may well prefer to see alternative outcomes!

Who do you call?

Ghostbusters?

But lets think about you and us as **users of the Internet**

These alternative options represent a pretty dismal future to the end user of:

- escalating cost,
- escalating application complexity and fragility
- massively reduced flexibility,
- stuff going wrong - massively increased risks of failure

Is there some economic factor at play here?

Right now individual short term business interests are leading the Internet towards collective long term suboptimal outcomes

At some point very soon the Internet will need some external impetus to restate short term interests to align with common longer term objectives

If we want IPv6 to happen we might need a large kick in the rear to get us there!

But what could be useful right now is ...



- An appreciation of the broader context of business imperatives and technology possibilities when confronting imminent IPv4 exhaustion
- An understanding that leaving things to the last millisecond may not be the wisest choice for anyone
- An appreciation IPv6 still represents the lowest risk option of all the potential futures

Fully deregulated environments do not necessarily make the wisest choices – this industry may need some additional applied impetus to avoid the failure option.

We may need a little regulatory push to see an IPv6 Internet emerge from this particular mess!

Tack

Thank you

Geoff Houston
gih@apnic.net

Patrik Fältström - paf@cisco.com