

Post IPv4 Exhaustion

Niall Murphy, Google
Dave Wilson, HEAnet

Health Warning

Central assumption of this talk:

**IPv6 doesn't
save us**

Structure

- Part 1: Cui bono?
- Part 2: A day in the life
- Part 3: Heaven and earth

Cui malo?

- IPv4 address allocation has uneven distribution
 - Very few have very many allocations
 - For example, 80% of space went to ~60 members in 2006
- Those who are most at risk from IPv4 exhaustion:
 - Those members who make the majority of requests?
- OTOH:
 - More likely to have addresses that can be harvested
 - More likely to have resources to deal with problem
- How about:
 - Those needing addresses without resources to cope if they don't arrive?

Cui malo?

- OTOH:
 - More likely to have addresses that can be harvested
 - More likely to have resources to deal with problem
 - Where do I want to spend global warming - The Netherlands?
- How about:
 - Those needing addresses without resources to cope if they don't arrive?
 - New entrants!

Cui bono?

- It's an ill wind that blows *nobody* any good
- Those who have an allocation who only ever needed a small amount of it
 - Business continuity
 - If/when markets emerge
- Those who can turn their more efficient use of IPv4 into a competitive advantage post-exhaustion
 - DSL providers vying to be the last to supply business customers with public statics
- IPv4 carrier-class NAT vendors
- ... ?

A Day In The Life

...life after the free pool

A day in the life

- Today typical network engineer day consists of:
 - Capacity planning
 - Addressing
 - Network design
 - Addressing
 - Adding new customers
 - Addressing
 - Monitoring, responding to incidents
 - Addressing
 - Reading random websites
 - Addressing

Monitoring

- Doesn't matter once everything is in site-local scope, right?
 - Depends... where did you plan to put the NAT?
 - Individual design decisions are "easy" but complexity means requirements will quickly begin to conflict
- Happy you can monitor through the NAT, wherever it lands?
- Happy NAT will do all the logging you need?
 - How is NAT Netflow looking these days, anyway?
 - NAT scales badly, but will monitoring scale it worse?
- Happy to increase acceptable failure thresholds?

Capacity Planning & Network Design

- My main job ("network planning") is reading minds
 - Turn "We need good connectivity to X" into "Z capacity to location Y, localpref W, on hardware P"
 - "With addresses A.B.C.D and E.F.G.H"
 - "Oh yeah, and J:K::L-M/N"
- How to answer these when a resource runs out?
 - Not good answers:
 - "But everyone knows..."
 - "Well you could just use IPv6..."
 - "We'll just use NAT - it'll only break some things..."
 - Only good answer:
 - no surprises, scare your boss early

Adding new customers

- Well, they'd better be single homed for a start
 - Lots of "hilarious" competition issues - more later
- NAT first, ask questions later
 - We do this already, so I guess move NAT one level up
 - Oh hey, weren't we shooting for single site-local scope?
- RFC1918 internally will cure all ills?
 - What breaks when you use RFC1918 on links?
(worse, what sort of works?)
 - IP multicast out the window, then?
(good job no one big around here uses that...)

Solutions available

- Still not clear to me that breaking the IPv6 chicken-and-egg cycle is inevitable, even when the chicken is in pain
- We could move quickly, at the expense of an unstable internet for some period of time, but either there is a trigger condition or there are workarounds - which?
- We have been given false hope by the ease of solving the client problem with NAT
 - What server-side solutions exist are far from mature (name two you'd deploy on your main website tonight)

Further layers of NAT

- A bit sucky if you're a user, or regular ISP, but will probably work for some definition of "work"
- Solutions if you are a colo provider:

...good job we won't be relying on a growing base of servers to support the growth of the internet, then.

Heaven and Earth

"There are more things in heaven and earth, Horatio,
than are dreamt of in your philosophy." - Hamlet

Heaven and Earth

So about that growing base of servers...

It's inevitable that in a NATted world, where the end to end principle is further and further eroded, application developers will have to shift further to a mindset where servers are "scarce"

New model of anything that can be reliably deployed goes from end-to-end and even pseudo end-to-end (e.g. skype supernodes) to a relatively small number of "superservers"

Heaven and earth

- Addressability and routability separate (mostly) now
 - This will change
- Allocations become valuable portals to DFZ
- Persistent TCP connections become more expensive economically
 - Require more kept state, but accrue less revenue
 - Ephemeral works
- Billing model switches to per-connection
- Incentives for designing apps moves towards server-centralised and ephemeral TCP

Billing changes... how?

Well, cost (if not pricing) is set by the bottleneck.

Right now that's bandwidth through the network.

- Link costs increase in rough proportion to bandwidth
- Routing costs increase in rough proportion to bandwidth and the size of the routing table

Billing changes... per connection

But NAT scales worse than both - it's CPU bound
(btw, you do know bandwidth use continues to increase faster than Moore's law, right?)

So not on day one, but ultimately, the cost of service provision will be transformed to *per-connection* and billing will go with it.

Some of us are better equipped to operate in such a world than others.