What now with IPv4?

This is a companion document to the TNC talk by Dave Wilson at https://tnc2011.terena.org/core/presentation/110

What is the current state?

IPv4 addresses have been managed by the Internet Assigned Numbers Authority (IANA) who maintained the free pool of unused addresses. These addresses were allocated to Regional Internet Registries (RIR) as needs required; the RIRs then subdivided these blocks and allocated them to networks as they required them.

On 3 February 2011, IANA reached the point where five blocks of addresses remained in the pool; these were given one to each RIR in accordance with previously agreed policy.[1] Each of the RIRs began to use the remainder of their addresses until reaching this final block, at which point special procedures kick in.

The Asia-Pacific regional registry (APNIC) has already entered this state. RIPE NCC, whose service region includes Europe, is likely to follow in the near future. It is very difficult to predict exactly when, as the impending runout may well result in a stampede of requests: see the the figure "APNIC Allocation Rate (smoothed)"[2]

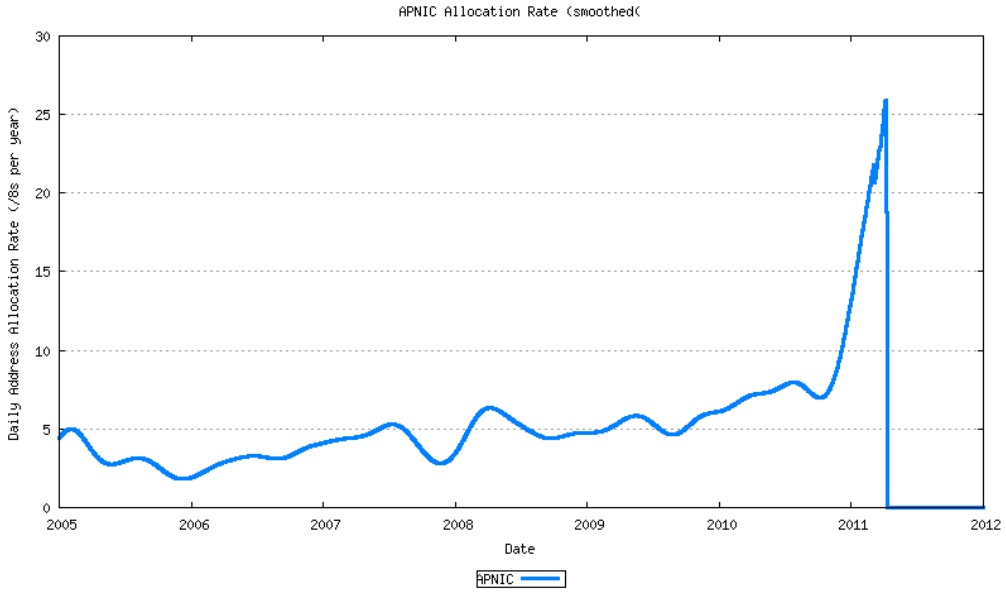

What are the immediate results of this?

Once the special procedures kick in here in the RIPE region, certain things are no longer possible:

- Growing institutions will be reliant on the remaining addresses available from their NREN.
- Provider Independent (PI) space, allowing connection to two networks at once, is gone.
- A research project that requires unfirewalled IPv4 connectivity may not be able to get it.
- Startups and new providers will only get a small, rationed portion of address space.[3]
- NRENs seeking to take new clients will not be able to request new addresses for them.

There are also impacts even in networks where IPv4 addresses are still relatively plentiful. All address assignments in the RIPE region - even those from an ISP to a customer - must now be justified with a plan to use most of them within six months. On 1 July 2011, this window drops permanently to three months.[4] Longer term projects must make repeated requests every few months, in the knowledge that no new addresses may be available at the time of the next request.

Transfers of address space are now permitted in certain circumstances.[5] This includes transferring for a cost. This not only means that we may have to spend money to get IPv4 addresses if other workarounds do not meet our needs; it also means that, in order to save money in the future, we may well be asked to do without our IPv4 addresses.

---

1 http://www.nro.net/news/ipv4-free-pool-depleted "Free Pool of IPv4 Address Space Depleted" - NRO

2 http://ripe62.ripe.net/archives/video/71 APNIC update to RIPE62, 4 May 2011

3 http://www.ripe.net/ripe/policies/proposals/2010-02 "Allocations from the last /8"

4 http://www.ripe.net/ripe/policies/proposals/2009-03 "Run Out Fairly"

5 http://www.ripe.net/ripe/policies/proposals/2007-08 "Enabling Methods for Reallocation of IPv4 Resources"

Workarounds are available, but many are not yet proven on the scales that they will be required. We already have a lot of experience with Network Address Translation (NAT) but until now it has always been possible, in extremis, to get more addresses. If we discovered that a firewall was particularly poor and unsuitable, or especially expensive given our requirements, we always had the possibility to get public addresses and remove it. With depletion, our only choice will be a different workaround, perhaps no less unreliable or expensive.

What about IPv6?

IPv6 is the long term solution, but it can't save us from the short term crunch.

In any given connection, IPv6 only works if your correspondent is also using IPv6, and the path between them works. For everyone else, we still need to have IPv4 available in some form or another, which means that deploying IPv6 does not save us from dealing with the problems caused by IPv4 depletion.

The original plan for IPv6 adoption was that IPv6 would be gradually rolled out across the internet, with traffic shifting from IPv4 to IPv6 as connectivity became available, with the expectation that this would complete in time for IPv4 run-out. This has not happened. So while we still need to deploy IPv6 in order to escape from the above problems on a permanent basis, workarounds are still needed in the meantime - and the IPv6 transition still needs help.

World IPv6 day - June 8th, 2011

On June 8th, 2011, a number of very high profile internet companies - including Google, Facebook and Yahoo! - will make their websites available to all users on IPv6.[6]

These organisations are already willing and able to roll out IPv6 access on their websites, and in many cases have been for some time. However, they face a problem that can only be solved by those of us on IPv4-only connections - some of our computers and networks are misconfigured.

Usually when a web browser encounters a site that is available on both IPv4 and IPv6, it will immediately detect if IPv6 is unavailable. However, a small number of users - a fraction of one per cent - are on machines or networks which do not fail over properly. This failure rate is too high for a large, popular website, so they are unable to proceed with their transition. However, without doing so, the problem goes undetected.

For a 24 hour period, beginning at 00:00 UTC on 8 June 2011, these websites and others will start serving IPv6 to all users who request it. If this experiment is successful (as others have been[7]) then some of these sites may choose to continue to serve IPv6 permanently.

IT departments may find it prudent to prepare for 8 June 2011 both by preparing their users and preparing their support staff for the event.[8]

Problematic machines and networks can be tested in advance by using tools such as test-ipv6.com - in order to detect misconfigured machines (as opposed to networks) this test must be run on the machine itself, so it may be wise to make users aware. Where a problem is found, it can be solved by removing incorrectly configured tunnels on the user's machine or, in the case of a misconfigured IPv6 network, restoring IPv6 access.

A strategic approach

After surviving IPv6 day, consider taking a strategic approach to IPv6 in future.[9] This is no longer a technical problem as such. The steps that HEAnet is taking to handle each of the problems are:

1. Decide on a strategic (not technical) level what action to take.
2. Calculate the impact on our IPv4 network and prepare for that impact.
3. Work on providing a full IPv6-only service, compatible with the IPv4 internet.
4. Assist our clients, providers and partners in managing the transition.

By seeing this as a strategic issue, rather than a technical problem, we can properly evaluate the costs both of action and inaction, and in so doing build a business case for the transition.

---

6 http://isoc.org/wp/worldipv6day/ "World IPv6 Day" - Internet Society

7 http://www.h-online.com/features/The-big-IPv6-experiment-1165042.html "The big IPv6 experiment" - Heise

8 http://ripe62.ripe.net/archives/video/168 "World IPv6 Day - Access Providers"

9 http://tnc2010.terena.org/schedule/sessions/show.php?sess_id=33 "A strategic approach to IPv6"