

IPv4 Address Utilization

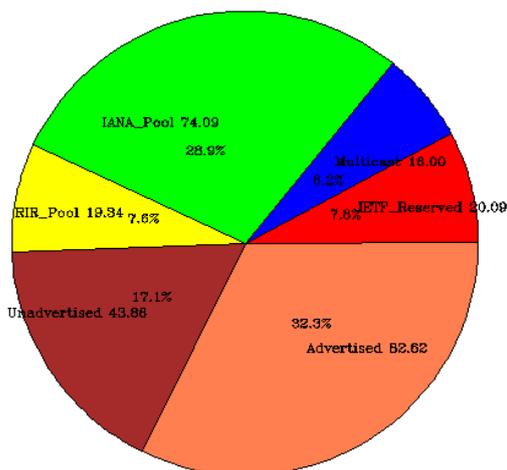
When the Internet Protocol was first designed some thirty years ago, one of the more innovative aspects of its design was the choice of an address field in the packet head that was 32 bits in length. At that time contemporary network protocols were using 8, or 16 bit address fields, embracing networks of between 256 and 65,356 separate hosts. The choice of a 32 bit field, embracing over 4.4 billion hosts was truly revolutionary, and at the time the address field appeared limitless.

However, in 2005 this is no longer that case, and while we will not run out of IPv4 addresses this month, and probably not in this year, it is time to look at how long we have to go, and what our options are when we exhaust the remaining available IPv4 address pools.

While the IPv4 space encompasses some 4.4 billion unique 32 bit values not every value can be used in the Internet. The IETF has reserved some 8% of the address space (or the equivalent of 20.09 of the 256 /8 network blocks). A further 6% (or 16 /8 blocks) are reserved for use in multicast contexts, leaving a little under 220 /8 blocks, or 86% of the total IPv4 address space available for general use.

The allocations of this space up until today now encompass some 66% of this useable space, or the equivalent of 146 /8 address blocks, leaving the equivalent of 74 /8 blocks in the unallocated address pool managed by the Internet Assigned Numbers Authority (IANA). So it looks like we're 2/3 of our way through the available address space, and asking the question of how long we have to go before we completely exhaust the address resource is a timely question.

There are some further aspects to consider here. The first is the procedure of address allocation. IPv4 address blocks are allocated from the IANA to the Regional Internet Registries (RIRs), who, in turn, allocate smaller blocks to ISPs and local Internet Registries, who in turn perform end user allocations. At this stage the address blocks are announced in the Internet's routing table. At any stage there are a certain number of addresses held in the RIR-managed address pools, and also a certain number of addresses held in the ISP and LIR pools before they appear in the routing table. Some 37% of the useable addresses are actually announced in the routing table, 9% are held in the RIR address pools and 20% of the useable address space is held in ISP and LIR address pools and is not announced in the Internet. Over one half of the RIR address pool space, and the majority of the ISP and LIR space reflects address space that was allocated prior to the adoption of the current RIR system in the late 1990's. Since 1997 over 95% of all RIR allocated address space is advertised into the routing system.



IETF Reserved	20.09 /8s
Multicast	16.00 /8s
IANA Pool	74.09 /8s
RIR Pool	19.19 /8s
Unadvertised	43.82 /8s
Advertised	82.82 /8s

Figure 1 – Ipv4 Address Space Snapshot - March 2005

By analysing the growth of the routing system we can derive a model of demand for address space across the global Internet, and, hence make some tentative predictions as to the longevity of the IPv4 address space. The current address consumption rate has lifted from a long term average of 4 /8 address blocks per year in 2002 to some 6 /8 address blocks per year in 2005. This implies that the remaining 74 /8 address blocks would provide a further 12 years supply at this rate, or until 2017 (Figure 2). Another form of the predictive model is where the growth of the Internet continues to increase over time, and the rate of increase of address consumption is 1.5 /8 address blocks per year. This model of continual increase in consumption will exhaust the available address space within some 7 years, or by 2012.

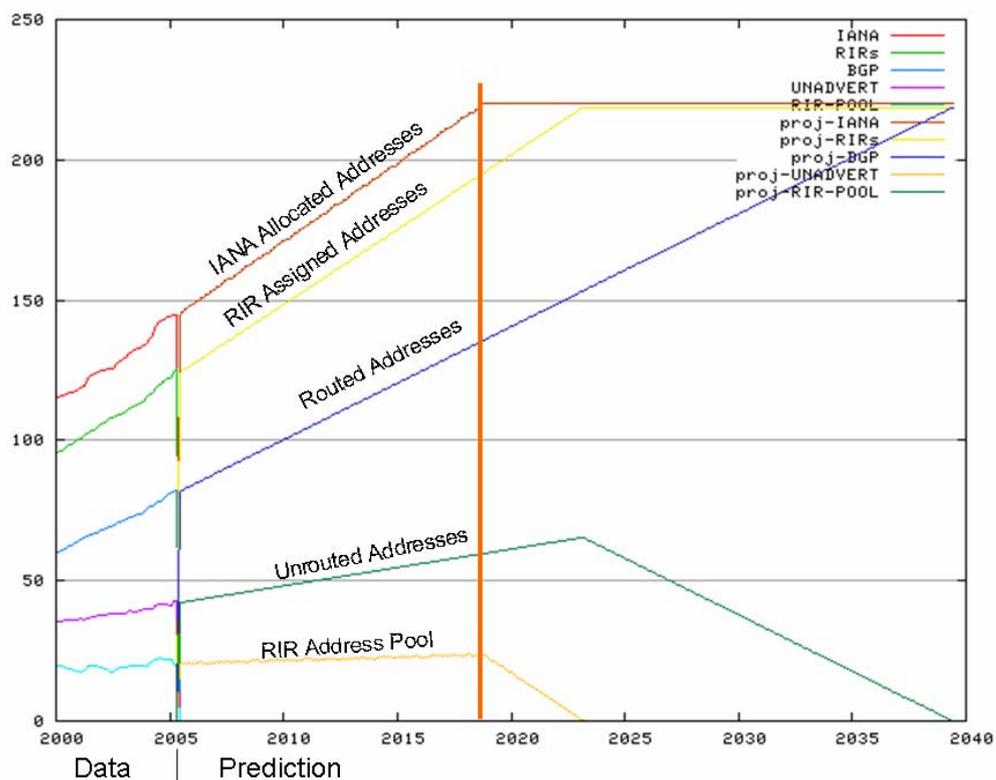


Figure 2 – IPv4 Address Consumption – constant rate model

These are approximate predictions, or course, and many technological, social and economic factors are at play when looking at the basis for actual address consumption rates.

The prospect of imminent exhaustion of the remaining unallocated address pools could fuel a run on the remaining space, dramatically increasing its consumption rates for the final few address blocks. On the other hand, the emergence of an address trading market could release a significant proportion of those unrouted address blocks that were allocated prior to 1995 into play, creating a new pool of address blocks that could fuel further Internet growth for up to a further two to three decades. Rather than relying of the emergence of a new market in address resources, or running the risk of a chaotic run on remaining unallocated address resources, it may be better to assist the Internet industry react to such preliminary signals of address shortage by undertaking an orderly transition to IPv6. Within such a transitional scenario the consumption rates of IPv4 addresses would slow down, with a corresponding increase in IPv6 allocation rates, as the Internet undertook such a protocol transition.

The basic message from this analysis is “Don’t Panic”. The Internet is not running out of available address space in the near future, and there is still quite some time available to work through the available options. The RIR system has been successful in ensuring that

address space is used responsibly, and there is no imminent exhaustion or shortage of IPv4 address space that would threaten the further orderly growth of the Internet at this point.

Of course this is a constantly evolving situation, and the policies used by each RIR to manage address resources is guided by an open, transparent bottom-up process that considers the perspectives of all interested parties when reaching consensus outcomes. In a realm where the finite nature of the resource demands careful and considered address management, the RIR system has proved not only to be an outstandingly successful approach so far, but also being easily up to the task of managing the demands of address distribution mechanisms in the coming years.

Geoff Huston
APNIC
March 2005