The Choice: IPv4 Exhaustion or Transition to IPv6

Jordi Palet
(jordi.palet@consulintel.es)

Introduction

- The IPv4 exhaustion is approaching and this may create some problems for the Internet, unless we make some smart choices
- These decisions need to be planned well in advance
- Waiting till the last minute will become very expensive, while early movers will have new and/or increased business opportunities
- This is going to affect the business of existing Internet Service Providers (ISPs) and to a greater extent, at a certain point in the time, the creation of new ISPs
- If your business uses Internet, it is going to be affected, sooner or later
- There are several potential ways to minimise the problems posed by IPv4 exhaustion

The Problem

- IPv4 is today the most widely spread version of the Internet Protocol
- However IPv4 is a limited resource
- IANA is responsible for the allocation of addresses to the RIRs. Each RIR serves a different region and distribute addresses to each of the five regions
- Current trends predict that the remaining addresses not yet been distributed by IANA are going to be exhausted around 2009
- If we do nothing, the existing Internet, in general, will keep working, but it will be quite challenging, indeed close to impossible, to keep it growing, unless we spend huge resources redesigning the applications and software embedded in all kind of devices
- Potential mitigations for the IPv4 exhaustion:
  - Temporary and Permanent mitigations

Temporary Mitigations

- Experimental IPv4 blocks
- Resource reclamation policies
- Changes in allocation policies
  - Global policies
  - Regional policies
- Secondary address market
- Increased usage of NAT

Experimental IPv4 blocks

- IETF reserved "experimental blocks", could be redefined and returned to the pool of "regular" addresses
  - several months (perhaps years) in terms of the standardisation effort, and afterwards it would require the vendors of hardware and software to release new versions, which could in turn mean further months or years. ISPs and users would then need to update their devices, which again, would probably mean several more months
  - not all existing devices will be able to be upgraded
  - it could be considered acceptable, in certain scenarios, for a "partial" upgrade to support those addressing blocks
- Time frame for this upgrade: minimum of three to four years; five to six years is probably more realistic on a global scale.
- Huge cost and it will only delay the exhaustion for a few months to a couple of years

Resource reclamation policies

- Before the RIR system was created, many IPv4 blocks were allocated to different entities ("legacy space")
- Since the creation of the RIRs, addresses have been allocated according to policies defined by the community through the Policy Development Process (PDP)
- One of the goals of those policies is the conservation of address space, in order to avoid wasting this scarce resource
- There are addresses that are not being used
- Even if we design a new routing protocol and implement it before IPv4 runs out, it may require renumbering
- Renumbering alone may be an alternative for using some of these addresses
- To reuse some of these addresses (that is, those that could be reused without new protocols and without renumbering), new policies could be created, typically developed at a regional level
  - to enforce the return of those addresses to the RIRs pools, but the community will find it difficult to actually exercise that action, and until that is done, the reclaimed address space will not be useful
- Much harder, to enforce the return of legacy space, as the legacy space holders aren't subject to the policies
- A couple of years' effort, implying some cost
- Some of the returned addressing space may be unusable (level of fragmentation)

Changes in global policies

- The communities of all the RIRs can use the PDP to create global policies. Global policies must be approved in exactly the same form (using the same text) in all the regions, and, once ratified by the ICANN Board of Directors, can enforce concrete actions to be taken by IANA
- This process could be used to create self-imposed restrictions on the community, for example in order to define a fixed date for the last allocation of IPv4 addresses to the RIRs, fix a reserve of them, or other actions that could be considered useful in preventing address exhaustion
- Challenging because of the political perceptions and/or implications regarding issues such as fairness of address distribution
- There is one more implication. If the community decides to impose restrictive measures in the subsequent use of IPv4, there will be legal implications, such as anti-trust considerations by governments and regulators in many regions. This, in turn, could affect the implementation of the policies and generate unpredictable problems with the RIR system and the PDP itself. It may even prompt governments and/or regulators to attempt taking over management of these resources
- Typically a couple of years from being drafted through to being ratified by ICANN and implemented
- Possible benefit in terms of how many years we are able to extend the address pool before reaching the exhaustion point is very limited compared to the effort and the risks

Changes in regional policies

- Changing aspects such as how allocations are done in the region, how utilisation of allocated addressing space is measured, and others. This might allow a region to more quickly obtain a bigger number of address blocks from IANA
- This might well be considered unfair by the other regions and perceived as a threat, which in turn could generate more aggressive policies to obtain more resources faster
- Not a desirable situation and it would quickly become a mess, generating a kind of “policy war” between the communities in different regions
- May take years and the benefit would not be significant in terms how much the exhaustion might be delayed

Changes in allocation policies

- One last important aspect of the PDP:
  - policies are meant to be changed at any time if it becomes convenient and the community agrees (for example, when new circumstances arise)
  - So policies (whether regional or global) that aim to impose self-restrictions in a permanent way may be misleading and actually conflict with the PDP itself:
    - the PDP cannot ensure that a policy, if approved, will not be changed again in a new cycle
    - so even if a policy is attempting to fix something for ever, such as usage of the last remaining address blocks, there is no way to avoid a new policy being proposed to cancel that immediately

Secondary address market

- Some organisations will find that it is economically feasible and beneficial to sell/buy IPv4 addresses (cheaper than the other alternatives). This might happen, at least initially, with legacy address space
- Current definitions of the system, addresses are not property and consequently selling them would be considered an illegal act
- Unless measures are taken to avoid it, this situation will create a “secondary address market”. This may damage the existing model, and to avoid this, it might be more useful to involve the community and the RIR system in defining how this market could behave in a way that it is coherent with the system
- It is difficult to predict how long it would take this address market to alleviate address exhaustion, and in any case, it is clear that it will become a very expensive solution
- It is also clear that acquisition of companies may become more and more relevant as a possible choice for obtaining addresses
- Existing work being done by the RIRs in what has been called “resource certification”, based on IETF standards, may be of use in preventing the growth of an address market and controlling legacy space
- If not adequately controlled by the community, will damage the allocation of resources by means of the existing RIR model and policy development process. This may turn the Internet into something totally different to what we have today, with possible implications for its neutrality
- It is expected that resource certification can be deployed in a couple of years or so

Increased usage of NAT

- NAT has been very useful, particularly in terms of connecting clients to servers. The combination of private addresses (which are also limited, but of which there are typically enough for even the bigger networks) and NAT allows the clients to start communications to servers, but does not allow servers to start communications to the clients
- Since the Internet started deploying not just client-to-server applications, but peer-to-peer services, NAT has become a nightmare, and actually has a lot of architectural implications for the Internet, reducing the possibility of end-to-end security and significantly increasing the cost of the development for peer-to-peer applications
- Enterprises and ISPs that are today using public IPv4 addresses for clients or even network infrastructure devices, may release those IPv4 addresses to be used by servers (which require them to be able to be accessed by clients with private addresses) somewhere else in their own networks. It may be the case that those ISPs or enterprises no longer need some or all of their public addresses, and these may then be returned to the pool of the corresponding RIR
- Doing this would have severe implications, not only in terms of the cost of renumbering the networks that are releasing those addresses. It may also be the case that small address blocks returned to the RIR may be unusable unless one or several contiguous blocks are available (meaning the resulting address block is big enough to be allocated to a new ISP)
- Operating NAT in ISP and enterprise networks typically entails additional support costs
- Even if it seems that increasing the deployment of NAT is feasible for client-to-server applications, there is a cost in renumbering the network and an associated cost in developing new peer-to-peer applications, especially if multiple levels of NAT are present. It is also possible though, that the release of IPv4 addresses associated with this operation (by renumbering to private address space) might go some way toward extending the life of IPv4

IPv6 transition (I)

- IPv6 has been designed with the principle of being able to coexist with IPv4 for a long period of time, avoiding breaking IPv4 networks and allowing all the existing services and applications to keep working without any disruption and allow a smooth transition from IPv4 to IPv6
- As more clients and servers support IPv6, more IPv6 traffic is being automatically generated, and consequently less IPv4 traffic. It is a very natural and transparent transition and in the very long term it could mean a phase-out of IPv4. There will be devices that can never be upgraded to IPv6; some IPv4 traffic will remain until all those devices disappear. There may come a point, tens of years from now, when we decide to simply ignore those IPv4-only devices and “force” their users to remove them or use “translators” that allow basic communication between IPv4-only and IPv6-only devices.
- If your devices are dual-stack, but your ISP connection, or part of it only supports IPv4, the most common situation, you need to use tunnelling mechanisms
- Upgrading ISP and enterprise networks to support IPv6 (actually dual-stack), should not, in general, present a huge cost, as long as it is adequately planned as part of a strategy for maintaining a network that is updated for new technologies, more bandwidth capacity, new services and functionalities, etc. Of course, this means that there are other related costs, such as training the engineers for the new protocol and operating a dual-stack network (instead of just IPv4)
- There are clear advantages to saving on the cost of supporting NAT-connected customers, and in a longer term, saving on the bandwidth associated with the use of tunnelling mechanisms if they aren’t done in the ISP network itself, but in third party networks
- Furthermore, supporting IPv6 instead of IPv4-with-NAT creates opportunities for new services and applications, especially peer-to-peer, which may be able to generate new revenue for the ISPs and in turn may also generate demand for more bandwidth
- Cellular networks already mandate IPv6 (supporting applications with IPv4 and NAT in that case becomes very expensive). So if interoperability between cellular and fixed networks is desired, IPv6 support needs to be provided in fixed networks

IPv6 transition (II)

- For end-users there is no cost in upgrading to IPv6. This means that even if the ISPs don't upgrade their networks (or only upgrade part of it) to support IPv6, the transition mechanisms will automatically allow users to connect to IPv6
- How quickly must this transition happen for it to be a useful solution to IPv4 exhaustion? Unfortunately, this is difficult to predict. We can estimate though that it would have to happen within a very short time, say 3-4 years. The big advantage of this transition is that it is not just a temporary solution. IPv6 address space is not unlimited, but it could be managed so that it will last us somewhere in the order of a few hundred years
- The cost is no higher than the others choices
- Some additional measures could be established:
  - Fees for IPv6 address blocks have been waived for a certain number of years in order to facilitate uptake of IPv6. In some regions this has not generated any increase in IPv6 adoption. In other regions though, especially those with less developed Internet industries, the fee waivers have been accompanied by other measures, such as dedicated IPv6 trainings and free support for setting-up IPv6 in ISPs, universities and public institutions. In these regions, positive uptake of IPv6 has been proportional to the level of training and support that has been provided, which indicates that those actions need to be continued or even extended
  - Further encourage ISPs' adoption of IPv6 through new regional policies. For example, there may be a way to attach new requests for IPv4 space to a simultaneous, mandatory request for IPv6 blocks (including, perhaps, some usage requirement). Additional criteria for obtaining IPv4 blocks could be imposed unless certain levels of IPv6 deployment are demonstrated (this might also make sense in the context of new IPv4 address blocks becoming increasingly scarce). Alternatively or in parallel, the cost of the remaining IPv4 blocks could be progressively increased over time, while simultaneously lowering IPv6 fees or extending any fee-waiving periods. Of course, increasing the yearly fees for those IPv4 blocks already allocated could also be helpful
  - New policies could be developed that make it more difficult to obtain IPv6 address space if the requester is holding unused IPv4 addresses. Fees can play also a part here, by relating the cost of IPv6 to the existence of unused IPv4 space: the more unused IPv4 space an ISP has, the more they must pay for IPv6. This might be a way to encourage an ISP to return some of its IPv4 blocks, even if it implies renumbering part of their network (the renumbering cost could actually be lower than the extra cost charged for IPv6 resources if the unused space is not returned)

Transition to a new protocol

- We could go back to IETF and develop a new protocol or a new mechanism to cope with the IPv4 exhaustion
- Our experience with IPv6 has shown that any new protocol can take several years to develop, several years to be implemented and several years to be deployed
- Even if we are able to invent a new way of re-using IPv4, we can expect that it will take a minimum 4-5 years
- The cost of this could be much higher than continuing with the IPv6 transition, it does not seem like a viable solution

IPv4 exhaustion phases

- Assuming we move towards IPv6 deployment:
  - IPv4 pre-exhaustion phase
  - IPv4 exhaustion critical phase
  - IPv4 post-exhaustion phase
  - IPv4 exhaustion very critical phase
- Timing discussed for the phases above assumes that routing scalability and/or routing protocols don’t improve over time and make possible other alternatives. Such alternatives might include being able to route every single IPv4 address, meaning fragmentation is no longer an issue and IPv4 addresses can be allocated/assigned one by one instead of by blocks

IPv4 pre-exhaustion phase

- Current situation
- Address scarcity is beginning to become apparent and there is an extensive usage of NAT, at least for residential customers
- IPv6 is not widely deployed, even if there are many transit networks offering dual-stack
- Some ISPs have deployed IPv6 in their core networks, just a few in their access networks and there is some IPv6 traffic, mainly using transition mechanisms
- We can predict that recent events such as the launch of Windows Vista, which comes with IPv6 enabled by default, will increase the usage of IPv6, at least in terms of residential end-users

IPv4 exhaustion critical phase

- When IANA runs out of IPv4 blocks and all remaining blocks have been allocated to the RIRs
- It may be that some new policies, possibly regional ones, will be adopted in order to increase the requirements and justifications for an ISP to obtain new IPv4 addresses
- There may be an important increase in the usage of NAT and private IPv4 space
- New ISPs may have difficulties obtaining as many IPv4 addresses as they wish from the RIRs; this will be much more visible in developing regions, unless some policies have been implemented in order to increase reserves at the RIR level
- Most of the big content providers already support IPv6 and as a consequence many old and new applications will make use of it
- Most of the ISPs will already have dual-stack networks, but this will not yet be widely deployed across access networks
- Will probably reach this point 3-4 years from now. It will last for a year or so, though this timing may depend on the availability of a secondary market (if this doesn't become useless due to resource certification), global policies and reclamation efforts. This period may also be extended if experimental blocks are reclassified by IETF and implemented by vendors (even given the interoperability problems that this may cause in some instances)
- The situation will most likely spark an explosive increase in IPv6 deployment on those networks that are lagging
- Cost of this deployment may be significant, but it may still be low compared to acquisition of IPv4 addresses on the secondary market, and with IPv6 traffic already becoming more prevalent, it would be worth the cost. There may be some networks starting to consider using only IPv6 (some may already be doing this in situations where IPv6 is already the dominant protocol)

IPv4 exhaustion very critical phase

- At this point the RIRs have also run out of usable IPv4 address space
- They may still have some resources left, but the fragmentation generated by resource reclamation may make it difficult to allocate these remaining resources under the existing policies
- It may become almost impossible to obtain more addresses from the RIRs, though this may no longer be an issue for existing ISPs, especially those that moved to a fully-fledged IPv6 network in the previous phase
- NAT usage is still increasing, but this is no longer a problem as almost every application and service now supports IPv6, with only a few client-to-server applications still IPv4-only
- Some sort of IPv4-to-IPv6 proxy or translation devices are deployed in some networks as part of the service provided to customers, especially on those networks deploying only IPv6 (this may create some interoperability problems)
- New ISPs are totally unable to get addresses from the RIRs, especially in developing regions, so they are unable to start businesses based on IPv4 with public addresses. However they may still use private addresses and NAT (even several levels), which should not be a problem as the reduced number of services available in IPv4 can be reached via IPv4-to-IPv6 proxy/translators boxes. Upstream providers would typically handle this as part of the transit service they provide. The upstream providers may also be able to assign downstream ISPs a small block of IPv4 public addresses, enough to make sure that the network works and no interoperability problems are created when connecting to other networks
- Very few ISPs or content providers are running only IPv4, and many access networks already have native IPv6 deployed
- This phase should encourage any lagging ISPs to move to IPv6: even if there is a cost involved, the business case for making the transition is now stronger than that for staying with IPv4
- This phase will be reached no more than one year after the end of the previous phase, and will last for a maximum of 1-2 years (though possibly much less if IPv6 deployment proceeds quickly on a global scale)
- More and more networks become IPv6-only, allowing the reduction of operation costs, while softwires is widely employed to support IPv4 traffic across those networks

IPv4 post-exhaustion phase

- Thanks to the massive use of IPv6 and sofwires, IPv4 addressing space within an ISP network is no longer needed; indeed, most of the IPv4 addresses being used for routers and other infrastructure devices may be returned to the relevant RIR pool
- This will facilitate new ISPs coming into the market without the problems described in the previous phase
- It may be that it is useful to renumber IPv4 in some networks, in order to return less fragmented IPv4 space to the RIR, however, the number of IPv4 addresses used in each network will be so low that renumbering will no longer be a hurdle
- It may be that the deployment of IPv6 becomes so widespread that it makes sense to consider dropping IPv4 from LANs – this is not recommended before this point in the time so as to avoid interoperability problems with IPv4-only services and applications. For this same reason, it seems clear that phasing-out IPv4 in the core and access networks may be a better approach that doing it first in the LANs.
- This phase may start at almost the same time as the previous one in some regions, depending on the degree of IPv6 deployment, and it will last for an undetermined period of time (until IPv4 is totally dropped, even from the LANs)
- It might be necessary to develop new policies, possibly global ones, in order to allow the transfer of IPv4 blocks between regions or back to IANA (so they can be allocated back, using existing policies, to other regions that may still require IPv4 address space). However, the level of global IPv6 deployment will probably make it unlikely that these sorts of policies will be necessary

The complete paper

- http://www.ipv6tf.org/index.php?page=news/newsroom&id=3004