The Choice:

IPv4 Exhaustion or Transition to IPv6

April 2007

Jordi Palet (jordi.palet@consulintel.es)

1. Introduction

The exhaustion of the remaining pool of unallocated IPv4 addresses is approaching within the next few years and this may create some problems for the Internet, unless we make some smart choices. These decisions need to be planned well in advance if we want to ensure that the changes are as trouble-free, painless and inexpensive as possible.

There is not too much time left; in fact the consumption of IPv4 addresses seems to be accelerating. If you're not following this situation and the possible implications, now is the time to pay attention: 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. As a consequence it may have a deeper impact in developing regions (Africa, Asia and Latin America/Caribbean) where the penetration of the Internet is not yet so widespread.

If your business uses Internet, it is going to be affected, sooner or later. Whether you act now or later will affect the impact on your business.

There are several potential ways to minimise the problems posed by IPv4 exhaustion, and it is becoming clear that several of these solutions will be adopted in parallel.

2. The Problem

IPv4 is today the most widely spread version of the Internet Protocol. This is the protocol used across the Internet for all kind of communications, and any device wishing to connect to the Internet needs to have an IP address.

However IPv4 is a limited resource. It is limited to a fixed number of connected devices: 4,294,967,296 (2^{32}, or 32 bits). This is the theoretical maximum, but in reality it is impossible to efficiently use all the possible addresses. It is also clear that even if we could use all of the addresses, they are insufficient to connect all the possible devices (and consequently people) that will soon require addresses.

The Internet Assigned Numbers Authority (IANA) is, for the time being, the organisation responsible for the allocation of addresses to the Regional Internet Registries (RIRs), which they do in blocks of 16,777,216. Each RIR serves a different region, and the RIRs, in turn, distribute addresses to each of the five regions: Africa (AfriNIC), Asia Pacific (APNIC), North America (ARIN), Latin America and the Caribbean (LACNIC), and Europe and the Middle East (RIPE NCC).

Current trends predict that the remaining addresses not yet been distributed by IANA are going to be exhausted around 2009. Many factors could affect this date in either direction, making it happen sooner or later, however, in recent years market evolution trends seem to indicate that exhaustion will most probably occur closer to the lower boundary of the predictions.

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 and gadgets.

We can classify the potential mitigations for the IPv4 exhaustion problem into two basic categories: Temporary and Permanent mitigations. These two categories are described in more detail in the following sections.

3. Temporary Mitigations for IPv4 Exhaustion

Let’s try to see a possible split of the different temporary mitigations. We can call them “Action Packages”, as each strategy may incorporate a set of different actions.

3.1. Experimental IPv4 blocks

The Internet Engineering Task Force (IETF) is the main body in charge of developing the protocols used in the Internet, including IPv4. The IETF originally reserved some blocks for purposes such as experimenting with new/future usages of IPv4. These experiments have not yet been done and are most probably not going to happen, however the reserved blocks don’t work in most of the devices connected to the Internet today.

In order to change this situation, those experimental blocks could be re-defined by IETF and returned to the pool of “regular” addresses. However, this will take at least 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.

Of course, not all existing devices will be able to be upgraded. In some cases vendors don’t exist anymore, so nobody is able to take care of the new software release, or the device may require a hardware change to be able to use the new software (which typically means being shipped to a service centre).

Nevertheless, it could be considered acceptable, in certain scenarios, for a “partial” upgrade to support those addressing blocks, instead of a global one. In this case, the expectation is that some networks may not be able to “talk” directly with others, making it difficult to predict communication problems when using those blocks across the Internet.

Moreover, it is clear that the time frame for this upgrade is a minimum of three to four years; five to six years is probably more realistic on a global scale. The cost of this will be huge and it will only delay the exhaustion for a few months to a couple of years. Of course, it may be the case that this “extension” will be enough for some service providers.

3.2. Resource reclamation policies

Before the RIR system was created, many IPv4 blocks (of different sizes) were allocated to different entities; this is the so-called "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.

Despite that, within both the RIR-allocated space and the legacy space, there are addresses that are not being used. Due to the fragmentation inherent in the design of the Internet (specifically routing), some of these addresses couldn’t be used at all, unless we design a completely new routing system (which in fact is needed, but as a longer term solution to cope with routing scalability problems). Even if we design a new routing protocol and implement it before IPv4 runs out, it may require renumbering, which, if even possible, would be a very costly task. Renumbering alone may be an alternative for using some of these addresses, however just try to imagine asking the ISPs, enterprises and even end-users to change the configuration of all their networks and connected devices. It would take many months or years and would have a huge economic impact.

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. These policies could 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.

It will be much harder, if not impossible, to enforce the return of legacy space, as the legacy space holders aren’t subject to the policies developed by the community, unless they voluntarily decide to bind to an RIR.

In any case, we are again talking about possibly a couple of years' effort, implying some cost in terms of human resources to actually do the job (most of the work could not be done in an automated way, as it requires ensuring that the resources are not used and then reclaiming them). Furthermore, only a few portions of the unused legacy space will actually be returned to the pool, for reasons including the impossibility of verifying whether the addresses are actually being used.

Moreover, some of the returned addressing space may be unusable; in many cases the level of fragmentation will make it useless.

Despite all this, and even if these measures provide only a small delay in the exhaustion of IPv4, it may be enough to allow some ISPs to keep moving ahead.

An alternative to “enforcing” this may also be considered and it is described below as part of Action Package 4.1.

3.3. Changes in allocation policies

The policy system is funded on the basis of policies that can evolve according to community wishes, based on consensus agreement. That means that the current allocation criteria could be changed in order to extend the availability of usable IPv4 addresses, and consequently delay their exhaustion. Two possible paths exist for this approach.

3.3.1. Global policies

The RIRs and IANA are bound by means of a Memorandum of Understanding between the NRO (Number Resource Organization, a coordinating and representative body for all the RIRs), the ASO (Address Supporting Organization) and the Internet Corporation for Assigned Names and Numbers (ICANN). As part of MoU, 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.

However, as indicated above, this will require reaching consensus in all the RIR service regions, which is usually a challenging task. In this case, we could expect that it would be even more challenging because of the political perceptions and/or implications regarding issues such as fairness of address distribution (among other possible issues).

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. Beyond legal issues, there are many other considerations that will need to be taken into account before moving ahead with self-restrictive global policies.

Last, but not least, if this global policy or set of policies is successfully developed, it will take typically a couple of years from being drafted through to being ratified by ICANN and implemented.

Once more the 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.

3.3.2. Regional policies

In addition to global policies, each region is able to develop their own policies. These policies could look at 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.

Obviously 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.

This is obviously not a desirable situation and it would quickly become a mess, generating a kind of “policy war” between the communities in different regions. Even so, this path may take years and the benefit would not be significant in terms how much the exhaustion might be delayed.

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). From that perspective, 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.

3.4. Secondary address market

There may come a point when some organisations will find that it is economically feasible and beneficial to sell/buy IPv4 addresses. Indeed, it may become cheaper than the other alternatives. It is easy to imagine that this might happen, at least initially, with legacy address space.

However, as per the current definitions of the system, addresses are not property and consequently selling them would be considered an illegal act. The organisations or individuals that were originally allocated the addresses are only “users” of the address space, which is effectively "leased" by the community.

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 (or parts thereof), which have legacy IPv4 addresses, may become more and more relevant as a possible choice for obtaining addresses. This could be considered part of this secondary market.

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. The basic idea is that only those resources registered with the RIRs will be certified, and that network operators could use this certification information to avoid routing non-certified resources across their networks, and consequently to the global Internet. However, it may be that owners of address blocks not certified within the system may be able to pay network operators to route those resources. It is difficult to predict whether we will reach that point and in what time frame, or if the initial model of “resource certification” will work without monetary interference for long enough to avoid that situation.

One point to be reinforced regarding the secondary market is that, if not adequately controlled by the community, it 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. It is also interesting to note that resource certification may be used to enforce Action Packages 3.2 and 3.3, as described above.

3.5. Increased usage of NAT

One of the protocols that has facilitated the widespread deployment of the Internet, despite the foreseen reduced number of IPv4 addresses, has been Network Address Translation (NAT). NAT has been very useful, particularly in terms of connecting clients to servers. It works by allowing many clients to connect to the Internet through a NAT device using a single public IPv4 address (or a set of them). The clients themselves use a special set of IPv4 addresses, which can’t be used in the Internet (so-called "private" IPv4 addresses). 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.

However it is important to remember that NAT remains a perfectly valid solution for client-to-server applications, and even if several levels of NAT devices are deployed in between the client and the server, it still works. Client-to-server applications are no more complicated (and no more expensive) when using NAT.

On the other hand, peer-to-peer applications are very expensive and complex to develop when even a single NAT is present.

NAT may be considered as an option for delaying the exhaustion of IPv4. 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.

However 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). This is basically the same fragmentation issue discussed in relation to Action Package 3.2.

Furthermore, operating NAT in ISP and enterprise networks typically entails additional support costs, as some applications don’t work automatically and it becomes necessary to manually configure some tricks in the NAT device (typically the access router). At a network operators meeting several years ago, an operator discussed the cost of NAT when taking account of the number of hours required to support those customers using it: “A single customer call to our first line of support costs us the revenue of that customer for a complete year; if the call needs to be directed to a second support line, it will mean the revenue of the entire life of that customer”.

As a consequence, 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.

4. Permanent Mitigations for IPv4 Exhaustion

We can foresee two Action Packages within this category.

4.1. IPv6 transition

IPv6 is a new version of the data networking protocol on which the Internet is based. The IETF developed the basic specifications during the 1990s. The primary motivation for the design and deployment of IPv6 was to expand the available addressing space of the Internet, thereby enabling billions of new devices (PDAs, cellular phones, appliances, cars, etc.), new users and 'always-on' technologies (xDSL, cable, Ethernet-to-the-home, fiber-to-the-home, Power Line Communications, etc.).

IPv6 has a 128-bit address space that can uniquely address 2^{128} network interfaces (actually 340,282,366,920,938,463,463,374,607,431,768,211,456 addresses, or about 340 sextillion).

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. At the same time, the way this coexistence works should allow a smooth transition from IPv4 to IPv6. In short, the basis of this coexistence and transition (not a migration, as it is not breaking existing networks and services/applications) is having both protocols in the hosts at the same time (this is called dual-stack), and allowing the operating systems and/or applications to choose which protocol they use for each communication (though in general the standards dictate that where possible IPv6 should be preferred over IPv4).

This means that 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. This last step, however, is very difficult to predict or complete, because, as noted in Action Package 3.1, there will be devices that can never be upgraded to IPv6; some IPv4 traffic will remain until all those devices disappear. Remember that having support for both IPv4 and IPv6 at the same time (dual-stack) means that even if some devices can only speak IPv4, they can still speak with those that already use IPv6. Of course, 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.

Now, what happens if your devices are dual-stack, but your ISP connection, or part of it (such as the access network, the last mile, or your ADSL link) only supports IPv4? This is currently the most common situation and, in fact, it was part of the design of this "transition plan" from IPv4 to IPv6. It is resolved by means of tunnelling mechanisms. A tunnel is a way to transport a protocol (IPv6 in this case) as data traffic across a network, which is based on another protocol (IPv4 in our case). There are many tunnelling protocols, some of which require manual configuration, but the most useful ones, from an end-user perspective, are automatic.

Dual-stack, tunnelling mechanisms and translators are part of a set of technologies that the IETF has called “transition technologies”, and together they form part of the transition plan from IPv4 to IPv6.

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). In general though, these costs will be minimal, unless the ISP decides to do it overnight. Today, all the enterprise and ISP routers, and all the operating systems already support IPv6 (in addition to IPv4), and this is generally true even for hardware purchased 3-4 years ago.

Of course, this scenario may be not true for very big networks that are outdated in terms of their equipment and software releases, and they may therefore require a longer period of planning in order to avoid the extra costs. There may also be other elements, such as very old equipment that will never be updated, that must be considered as part of the cost, as some translation or other complementary technologies may be required.

Given that the cost of upgrading to IPv6 is so insignificant (across a network update cycle, which may mean months or a few years, depending on the specific case), there are clear advantages to saving on the cost of supporting NAT-connected customers (as in our example in Action Package 3.5), 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.

On the other hand, 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. Cellular phones have had IPv6 support for the past 2-3 years, which typically means every user today has an IPv6-capable cellular phone.

For end-users there is no cost in upgrading to IPv6 for PCs and similar devices, and many other devices already support it with simple upgrades (some networked-printer vendors, for instance). The same is true for many other devices, such as IP cameras, videoconferencing equipment, gadgets, etc. In fact, over the past 4-5 years virtually all operating systems have received some sort of IPv6 support.

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.

The question then is 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, if it was to compete with the other Action Packages in terms of cost. 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.

None of the other solutions discussed here offer this, and the transition to IPv6 also offers the positive effect of restoring the Internet's end-to-end principles, thereby enabling innovation opportunities, which in turn will provide new revenue opportunities.

Last, but not least, the cost of this Action Package is no higher than the others, and could possibly help to reduce some of the existing costs as well.

As part of this Action Package, some additional measures could be established. For example, in some RIR service regions, 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.

Another possible alternative is to further encourage ISPs' adoption of IPv6 through new regional policies. Those policies could act on many different fronts, of course, and may not produce the same results in different regions. 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.

As an alternative to Action Package 3.3, especially for the holders of legacy IPv4 address space, 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).

4.2. 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. As explained earlier, this was the main goal when developing IPv6. As our experience with IPv6 has shown, any new protocol can take several years to develop, then 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. Considering the cost of this could be much higher than continuing with the IPv6 transition, it does not seem like a viable solution.

5. The IPv6 Deployment

Let's assume that we move towards IPv6. In reality, this is already happening and some facts are pushing it now more than ever before. It is mainly occurring on two different fronts:

a) **Availability of IPv6 in the core of ISP networks.** In general, many backbone/core/distribution networks of ISPs have already made the move to support IPv6, either natively (without the need for transition mechanisms) or with some kind of tunnels (including technologies such as MPLS). In many cases, this was planned long ago, and as explained already, this has typically meant that the cost has not been a problem. There are others who already have concrete plans for supporting IPv6, and it is a matter of the appropriate timing for each case.

b) **Availability of IPv6 in the end-user platforms.** As discussed earlier, many operating systems have supported IPv6 for some years now, with no additional cost. It is fair to say that all operating systems being marketed today support IPv6. Only a few outdated operating systems, such as older members of the Windows family (95, 98, ME) lack IPv6 support.

So, what makes the difference now? As time passes, more and more core networks are receiving IPv6 support – in a few of years it will be inconceivable that any medium/big ISP would not be IPv6-ready. In fact there are already some business arguments in favour of moving to IPv6, such as public procurements that already ask for IPv6; not losing customers is a big enough reason for shifting. This will probably happen more often as time goes on, in a kind of domino effect: more public networks become IPv6-ready, forcing others to do the same.

On the other hand, Windows XP, the operating system that has the biggest market penetration and has supported IPv6 since its release, is being replaced by a new version, Windows Vista. The main difference between them, from an IPv6 perspective, is that in XP IPv6 must be enabled either by the user or the applications (some already do so), while in Vista, IPv6 is enabled by default and the built-in operating system services use it (by means of automatic transition mechanisms, so it works even if the ISP doesn't offer IPv6 support). New versions of applications such as Windows Office also make extensive and preferred usage of IPv6. Peer-to-peer applications work better and faster with IPv6, may require less bandwidth and can be up to 80% cheaper to develop than IPv4 applications (because of the NAT implications), which also means a much shorter development time, and other application vendors will likely start taking advantage of that.

In fact, since the release of Vista (November 2006), it is already possible to measure an increase in IPv6 traffic. This does not mean end-to-end native IPv6 traffic, because the majority of edge/access networks (the last mile to reach the end-user) don't yet support native IPv6. The automatic transition mechanisms embedded in the operating system, however, allow Vista to use IPv6.

Whether we like it or not, the snowball has been launched. Windows' market share is possibly around 75-80% and it is reasonable to assume that a huge majority of end-users will have upgraded to Windows Vista within the next year. Other operating systems, such as Mac OS X and other open-source alternatives, already have IPv6 support enabled by default. As a result of all this, most peer-to-peer traffic may soon be across IPv6, and peer-to peer traffic already accounts for the biggest portion of total Internet traffic. It is not unreasonable to predict that within two years up to 50% of Internet traffic may be already IPv6.

The operating system update cycle in larger enterprises is typically a bit longer. Many of them, unlike end-users, may be still running non-IPv6-ready operating systems. However the availability of applications using IPv6 will also make a difference and even when using XP, business customers will want IPv6 enabled in more and more cases.

Many enterprises may initially turn-off IPv6 in Vista, whether through security concerns, lack of knowledge (which may mean lack of IPv6 support to users), or lack of support for dual-stack in their network. The last case should not be a problem, as Vista is able to use IPv6 on an IPv4-only network via transition mechanisms, but clearly this, together with a lack of knowledge, may well delay the deployment of IPv6 in many enterprises. This will change over time and possibly when an enterprise updates its firewalls with new software releases, this could also include an upgrade to support IPv6.

What is missing then? The biggest offenders are the access networks and the content providers. How important is that? For the time being, not very.

As mentioned earlier, IPv4-only access networks can be traversed using transition mechanisms. Furthermore, it is understandable that ISPs don't yet have any interest in updating their access networks, because there is no economic incentive to do so and no new services that would help to pay for the upgrade cost. Most of the low-cost residential routers (Customer Premises Equipment, or CPE), which are typically provided by the ISP, aren't yet IPv6-ready. ISPs will only replace those CPEs when there are other reasons for doing so, such as providing new technologies, and/or offering more bandwidth. Of course, there will be cases where those devices could be upgraded with new firmware versions, and this may actually be encouraged by ISPs, but typically it is only the more experienced users will do this (on-line gamers, for example).

Regarding content providers, there is a similar situation. Right now, there is no incentive to move to IPv6. No new revenues are foreseen, at least not until new applications can be offered that take advantage of IPv6. On the other hand, for the most common usages of the Internet (browsing, emailing and client-to-server applications or services), there is generally no practical advantage in using IPv6, because they work fine even if the user is behind one or several levels of NAT. New multimedia-based applications, however, which work peer-to-peer, may find a better “quality of service” (QoS) over IPv6 than they do over the current IPv4.

When a medium/big content provider in a given country or region starts offering an IPv6-only service (or even content available only in IPv6), other providers in that region will realise the business disadvantage of not moving to IPv6. Obviously this will get much more interesting when a single global content provider starts offering IPv6, as will hopefully happen sooner rather than later.

Moreover, the cores of most of the bigger networks, including the most important multinational carriers, are already IPv6 enabled.

In regions such as North America, networks are generally less IPv6-ready than in other regions. For example in Japan, and to a lesser degree Europe, some ISPs already provide IPv6 up to the edge for residential customers; this has not yet happened in North America. Similarly, a much larger percentage of ISPs in Japan, Europe and even Latin America, support IPv6 in the core of their networks than in North America.

Most of the research and educations networks, and many universities in Japan and Europe already have support for IPv6, but not that many in other areas. In contrast, there is almost no visible IPv6 adoption in the business world, with the exception of a very few sectors and a couple of multinational corporations, which have already seen the opportunity for some specific applications and services.

So, considering that the transition to IPv6 is already happening, should we try to be more aggressive with IPv6 deployment? This is something that ISPs should seriously consider, because as more and more IPv6 traffic traverses their networks by means of transition mechanisms, there will be an impact on the perceived QoS that users get from their network, as well as additional upstream provider bandwidth consumption.

This is especially important in developing regions, where delays are already an issue (because of lower network capabilities, usage of satellite, etc.), and the upstream bandwidth cost is terribly expensive.

Why does this happen? In short (and without entering into very technical details), we can say that the automatic transition mechanisms, at least the more important ones such as 6to4 and Teredo (which works by traversing NAT), can be very efficient within a single provider network or when used across different providers that exchange traffic. They are also very efficient when different clients talk over IPv6 using the same transition mechanism (i.e. 6to4-to-6to4 or Teredo-to-Teredo).

However, as much more clients start using IPv6, the chances increase of users employing different clients, even in the same ISP network, and using different transition mechanisms (or a transition mechanism and native IPv6). That implies inserting a third party relay into the path of IPv6 traffic, even between two clients connected to the same ISP. This, obviously, means extra delays and extra upstream bandwidth. For this reason, there is a lot of interest in deploying those relays within the ISP networks, even if that ISP has not done a complete IPv6 deployment, or does not have any native IPv6, because it would help to reduce costs and avoid the perception of a degraded QoS. The cost of deploying boxes with those "relaying" functions in the networks is meaningless, as routers already can do this (as can any PC running Linux, BSD, Windows or other operating systems) without any special operation or maintenance required. Moreover, one could take measurements regarding how much traffic is being relayed on those boxes - this data might act as an incentive to move more quickly to a wider native IPv6 deployment.

One way to be more proactive with the deployment of IPv6 in a given network is to use a tunnelling protocol that works automatically but at the same time provides end-users with IPv6 addresses that are from the ISP's own address pool, instead of IPv6 addresses provided by the transition mechanism itself (as happens with 6to4 and Teredo). Softwires is an IETF protocol, basically L2TP (Layer 2 Tunnelling Protocol), which offers an alternative in this regard. Softwires allows for the automated creation of tunnels in order to transport IPv6 across an IPv4-only access network or even to an IPv4-only core. Furthermore, softwires can also be used in reverse, when you need to transport IPv4 traffic across an IPv6-only network.

So what can ISPs do to help accelerate IPv6 deployment so that it works in their favour instead of against them? It can be summarised as follows:

a) **IPv6 prefix.** For many years now, all the RIRs have been allocating IPv6 address space. It should not be an issue for you, as an ISP, to calculate how big your network is, how much you expect to grow and submit an appropriate request.

b) **IPv6 transit.** Be sure to obtain IPv6 connectivity from your upstream provider/s, typically the same ones you use for IPv4 (so yes, we are talking about dual-stack). Obviously, if native IPv6 (dual-stack) is not possible, tunnels will do the job. If your upstream is unable to provide IPv6 connectivity, you should probably follow the chain upstream and obtain connectivity from the upstream of your upstream, or even a third party service provider.

c) **IPv6 transition support.** If part of your network, typically the access, can't be upgraded to support native IPv6, then you need to keep as much of the IPv6 transition traffic as possible within your own network. This can be done inexpensively by means of configuring relays (6to4, Teredo). At the same time, you can use other transition technologies, such as softwires (described above).

d) **IPv6 peering at Internet Exchanges (IX, IXP, NAP).** This is that the same as for IPv4: it is preferable to avoid using upstream bandwidth for local or regional traffic exchange.

e) **Native IPv6 core.** It is important to plan for the upgrade of the complete network, but in most cases it is easier to upgrade your core network first, and upgrade the access network at a later stage. This should not be a problem, especially if you took the steps outlined above in point b).

f) **Native IPv6 access.** Typically, your access network will be a bit more resistant to upgrading, depending on the technology being used. However, it is simply a matter of planning. Sooner or later, you will be offering new features that will require that you replace the CPEs in your access network, and then you will also be able to provide native IPv6.

g) **"Reverse" transition.** At a certain point, your IPv6 traffic may overtake your IPv4 traffic. This is an appropriate time to consider starting to use softwires and removing IPv4 from your core network (possibly also from the access network). The idea is to keep dual-stack (even using NAT with IPv4 private addresses) in all the local networks (including data centres, which in this case should still use public IPv4 addresses), so you can avoid translation of protocols that may not "survive" translation intact, and automatically encapsulate IPv4 in IPv6 from the local area networks (typically should be done by the CPEs, but it may also be done by the hosts).

h) **IPv4-to-IPv6 translation.** Translation is not perfect, as IPv4 and IPv6 are non-compatible protocols, but applications that work with NAT, typically client-to-server applications, using "simple" protocols (for example, http), work just fine with proxies or translator-like devices. At some point in time, there may be devices which are IPv6-only (though this is not recommended for the time being, especially when you can avoid it by using private IPv4 addresses behind NAT together with IPv6) and these may need to talk with old boxes which are IPv4-only. Translation/IPv4-to-IPv6 proxy may also help you to reach a point where IPv6 traffic is dominant over IPv4 within your own network. It can also mean that even if you still carry a lot of IPv4 traffic (typically http) to the rest of the Internet, you can reach the status described above in point g), and carry only IPv6 natively.

These steps are not necessarily to be taken one by one or in the order in which they are listed above. In fact, typically steps a), b), c) and d), can be done simultaneously, or very close in the time, without too many difficulties. Steps e) and f) will depend much more on the exact situation of each network and any available upgrade plans or maintenance cycles. In many cases, step g) could be also part of c), e) or f), and typically when you start with g), it means you no longer need c). Similarly, h) may be required together with g), or even before if your network becomes one of the cases of IPv6-only networks, which are already happening in some sectors and may become more frequent as we move toward a global IPv6 Internet.

6. IPv4 Exhaustion Phases

Let’s return to the question of IPv4 exhaustion. We know it is coming and we have some idea about when it will happen. However, IPv4 is not something that we can turn completely on or off. This transition is most likely to occur in several phases, so let’s try to classify those phases. Let’s also assume that we agree that the deployment and use of IPv6 is increasing.

6.1. IPv4 pre-exhaustion phase

This is the 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.

6.2. IPv4 exhaustion critical phase

This describes the situation 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.

We 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 (3.1) and reclamation efforts (2). This period may also be extended if experimental blocks (1) 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. The 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).

6.3. 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. As a consequence, 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.

At this stage 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.

6.4. 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 (Local Area Networks) – 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.

Note that the 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.

7. The Smart Choice

It is clear that IPv6 is the right choice. This document presented a view of all the alternatives or Action Packages, and even if many of these happen in parallel, IPv6 provides the only long term and stable solution.

There are many examples of this transition happening already and many others that demonstrated why this is a necessity from other perspectives. We could just mention DOCSIS (Data Over Cable Service Interface Specification), which is the protocol used to carry data in cable networks. The new version of this protocol (v3.0) adopted IPv6 in order to allow for bigger cable networks to be deployed. Previously cable networks used private IPv4 addresses, but those blocks are already too short to keep growing and match the continued demand for broadband deployment. There is little doubt that in the short to medium term, there will arise other situations for which IPv6 offers the only solution.

Yes, some people say IPv6 is basically IPv4 with 96 more bits, and on a basic level, this is true. However those extra bits make a huge difference if we really want to maintain a stable Internet over the medium to long term: in this case, "size matters".

Coexistence and transition are necessary steps to avoid disruption of the Internet, and as in the case of any major technological change, it may involve efforts and costs, even pain, but this will be much less of a problem than those we will face in the longer term if we don't make the appropriate choice. The advantage is the future of the Internet itself. It is a matter of striking a balance between investing for the future and spending today.

A global IPv6 Internet is what the future deserves, not an immediate change, but a smooth one where IPv4 will be present during many years, perhaps even forever. There will come a time though, when IPv4 traffic will be insignificant.

There are very important additional issues that need to be considered, such as the level of innovation that the deployment of IPv6 could bring back to the Internet. This is especially relevant in developing regions, where a continued focus on using IPv4 will hinder further development. IPv6 needs to be put in context as an innovation opportunity with a value that IPv4 can't provide by any means.

8. Conclusions

From all that has been introduced in this document, it is clear that the smart choice for a long-term solution to IPv4 exhaustion is IPv6. The key is to begin planning for IPv6 now.

Good decision-making at this key point will help us to bridge the digital divide; the wrong decisions will serve to increase the gap. Governments, regulators and other public bodies have an important part to play in guiding this choice. We must not advocate for enforcing pre-determined decisions onto the market, but instead facilitate the community in making the best choice, ensuring that there is sufficient awareness of all the affected stakeholders and that public tenders demand IPv6 support, preventing public resources being wasted when IPv6 support becomes a necessity.

A final issue to consider is the question of fairness in the distribution of IPv4 addresses. This argument is commonly used to propose new global or regional policies, but it is absolutely unrealistic. The Internet and its evolution are constantly changing. The distribution of addresses based on factors such as population, deployment levels, services, applications or other factors, are subject to changes (changes in population levels, movements across regions, etc.). Basing policy on factors that are likely to change can never be considered fair, because those same policies may become unfair as the original factors, or any others which are part of the formula, change.

The fairest distribution is to ensure global availability of resources and this is only possible with “more bits”: IPv6.

Last, but not least, IPv4 exhaustion is not just a problem for ISPs or Internet users. It is a global issue, but software, application and service developers (among other industries) have a big stake in this and could take advantage of that. Developing anything using IPv6 instead of using IPv4-with-NAT is simpler, cheaper and may be much more powerful. The lower cost can be invested in more features, and this is to the benefit of everyone, for the better development of the Information Society.

Have you already made your choice? If not, be aware that some already have, and you are beginning to be at a disadvantage.

9. Acknowledgements

This paper is partially inspired by numerous conversations with many community members in all the RIR service regions, in addition to long discussion on numerous mailing lists (mainly those related to RIR policy).

I would also like to acknowledge the inputs received from Leo Bicknell, Chris Buckridge, Roque Gagliano, Paul Rendek, Paul Wilson and Bill Woodcock.