

## Internet Addressing in the 2010s

### IPv4 exhaustion and address transfers, and their impact on IPv6 deployment

Paul Wilson, APNIC



#### Important notes:

1. For readers requiring background information on IPv4 and IPv6, it is recommended to read Part 2 first (page 7).
2. Throughout this paper, "IP" refers to the Internet Protocol, not to Intellectual Property.

## The IPv6 Challenge and Opportunity

It seems well known these days that the Internet is running out of the address space that is required for future growth. The current address space supply, provided by the IPv4 standard, is close to exhaustion; as a result, sustainable future growth will require the deployment of IPv6, the "next generation" Internet Protocol.

IPv6 was standardized in 1998<sup>1</sup>, and implementations have been available for many years. However its practical deployment and use on the Internet have been much slower to eventuate than many expected. The primary reasons for this are not technical, but related to economic and business factors in Internet services provision.

1. <http://tools.ietf.org/html/rfc2460>

It has often been observed that IPv6 deployment comes without direct incentives to encourage its adoption, but rather brings several “vicious circles” of supply and demand dependency:

- IPv6 is designed to invisibly replace IPv4, making no immediate difference to Internet users; therefore Internet Service Providers do not experience demand for IPv6 from their customers.
- Consequently ISPs and online content providers have tended not to prioritize IPv6 service development, or to demand IPv6 support from their own suppliers.
- The cost of IPv6 deployment is significant, and higher returns have been available from investment in other service improvements, such as access network infrastructure and bandwidth.
- The once-hoped-for “killer application” for IPv6 will not be developed until there is an infrastructure that it uniquely requires, but neither can these imagined innovative services encourage IPv6 deployment until they actually exist.

IPv6 has also suffered from a range of disincentives for “early-adopters”: a more limited range of products and services supporting IPv6, higher risks and costs in

use of those products and services, and the need to develop technical expertise. Many Internet business relationships (such as network interconnection, content and application distribution, and client-server delivery models) rely on the use of a common protocol by all parties, so early adopters cannot benefit from IPv6 until others follow.

These costs and challenges should not be overestimated; indeed in some cases they are quite marginal. However they exist in an Internet business environment that is highly competitive, where costs and capital investments must be carefully linked to returns, and business planning tends to be short-term. All of this means that the IPv6 business case must be explored and proven, and even *that* process requires new management information and expertise.

The economic reality and dynamics of IPv6 deployment reveal fairly clearly why a widespread transition to IPv6 has not yet happened as anticipated; and why many have advocated that external incentives may be required (see “Actions for Government” below). On the brighter side, it may be expected that once a point of “critical mass” is reached, natural incentives will take over and complete the process, however this point still appears to be far off.

## IPv4 is Dead: Long Live IPv6?

In the absence of demand for IPv6 itself, the critical driver towards deployment can only be the lack of IPv4 address space required for new Internet infrastructure. In this sense, the “killer application” for IPv6 is the Internet itself, and not any specific application, product or service that IPv6 can enable.

IPv6 might therefore seem to be an inevitable outcome of IPv4 address exhaustion; one which will happen naturally and as required, without intervention. This is only partially true, for a number of reasons. In theory, the vicious cycles described above should be broken when IPv4 exhaustion “arrives” and all participants are compelled to move to IPv6. In practice, however, the motivation towards IPv6 will never be so evenly distributed or so predictable that such a day could come. This is due not only to the distributed and uneven nature of IPv4 consumption, but also to techniques that are available to extend the lifetime of IPv4, which are now discussed.

### IPv4 Lives On: NAT

A major contributor to the longevity of IPv4 has been the deployment over many years of Network Address Translation, otherwise known as “NAT”. This is a widely-used Internet technology that can greatly extend the practical use of a single IP address, by sharing it among by many devices. Such devices exist on a so-called “private network”, which is located “behind” its single public IP address. Each individual device receives a distinct private address (from a designated “private address” block<sup>2</sup>); however, to the

rest of the Internet the entire collection of so-called “NAT-ed” devices appears to be located at one address. As only one public address is used by a potentially large collection of connected devices, NAT deployment has already helped to avoid the consumption of many IPv4 addresses. In fact, without NAT, IPv4 exhaustion would have happened many years earlier than it has.

The function of a NAT on the Internet can be viewed much like that of a PBX (Private Business Exchange) in telephony. The telephones that sit behind the PBX can access the global telephone network, and they notice little difference, though there are certain penalties in efficiency and functionality of their operation. From the outside of the PBX, for instance, it is actually impossible to reach a given “private” telephone number (extension) without going through an extra step in the “protocol”, normally a conversation with a human or automatic operator.

While PBXs bring some disadvantage to telephone users, NATs bring much greater penalties to the Internet, due to its need for fast response, reliability, automation, and standardization of communications transactions. In small-scale use (very commonly in the case of home- or office-based Internet gateway devices) this penalty may go unnoticed by the user, because applications and service providers have invested the efforts required to ensure their products are compatible with NATs of all kinds.

---

2. <http://tools.ietf.org/html/rfc1918>

However, at large scales, the use of NAT in providing Internet services to thousands or millions of users will incur much larger penalties. So-called CGN (Carrier Grade NAT) devices are still not widely available, but are being developed in response to IPv4 address space shortage, to provide one avenue for expansion of access networks without the use of IPv6. Such devices will be expensive, yet still limited in their capability to support arbitrary Internet traffic and services, particularly services which are not yet in common use.

In the long term in fact, widespread use of NAT throughout the Internet would ultimately destroy the current “end-to-end” model of the Internet, under which any device can potentially exchange arbitrary traffic with any other.

## IPv4 Lives On: In an IP Address Market

Before the exhaustion of IPv4 address space, there was little incentive for addresses to be transferred between one party and another. Those who needed addresses were able to receive them readily from their appropriate IP address registry (RIR), at reasonable cost and by following well known procedures. While there may have been some “black market” address transfer activity, this appears to have been very limited, and mostly confined to those who wished to conceal their address usage (e.g. to conceal hacking or spamming activities), rather than as an alternative means for obtaining addresses.

As IPv4 exhaustion approached, however, it became clear to many that a transfer mechanism could be useful in at least three ways:

- (1) Establish a means for those needing additional IPv4 addresses to receive them
- (2) Provide holders of unused or under-utilized address blocks with an incentive to release them for use by others
- (3) Reduce the incentive for a black market to emerge.

A contentious issue in discussions about IPv4 transfers was that of whether the principle of “demonstrated need” should be applied; in other words whether an RIR, in registering a transfer, would require the recipient to demonstrate their need for that address space (as they had to previously, when requesting blocks from the RIR’s free pool). The strong argument in favour of a “demonstrated need” requirement was to avoid the emergence of IPv4 stockpiling and artificial scarcity; while the argument against was to eliminate any barriers that may encourage black market transfer activity.

At this stage (April 2013) two RIRs – APNIC and ARIN – have active address transfer policies, both including a “demonstrated need” requirement<sup>3</sup>. Additionally these RIRs permit inter-regional transfers to take place between their regions, and in future with other RIRs which institute compatible transfer policies.

3. <http://www.apnic.net/policy/transfer-policy> and <http://www.arin.net/policy/nrpm.html>

## Transfer of Legacy IPv4 Space

It should be noted that with respect to IPv4 transfers, neither APNIC nor ARIN make any material distinction between “legacy” and “current” IPv4 address space: blocks in either category can be transferred both within and between these RIRs’ regions.

There is a common perception that the issues of IPv4 transfer and legacy address space are connected in some way, however there is no relationship in terms of transfer policies or mechanisms. It is likely that legacy address blocks will be transferred in greater volume, because a larger proportion of this space may be under-utilized, and thus more easily released for transfer. This is in fact one strong benefit of a transfer mechanism, because without it there has been little incentive for legacy address space to be used more efficiently.

## IPv4 Transfers and IPv6 Transition

It has been suggested that the availability of IPv4 transfers will further reduce the urgency of IPv6 deployment and delay the transition process. This is likely true, because as mentioned above the biggest driver for IPv6 deployment is the depletion of IPv4 address space.

That said, the limited incentives for IPv6 transition create a requirement for concurrent action by as many parties as possible, and the current uneven distribution of IPv4 address space (particularly considering legacy IPv4 stocks in North America and Europe) defeats this requirement. Therefore while IPv4 transfers

may indeed reduce the urgency for some parties, they will also serve to improve the distribution of “motivation” towards IPv6 in the longer term, which may be essential to a successful transition.

In any case, any delaying effect created by IPv4 transfers will be strictly limited by the relatively small supply of IPv4 addresses that are practically available for transfer. Assuming that 10% of the total IPv4 address space was readily available for market transfers, then at 2011 consumption rates, this additional supply would last less than two years.

It is clear that the interactions of the emerging IPv4 market and IPv6 transition are complex and untested, and merit further investigation. The situation is attracting much interest, and it has been one aim of this paper to clarify factual aspects, while shedding light on those dynamics.

## IPv4 Lives On: In a “Dual-Stack” Internet

The transition to IPv6 is not a single event, either globally or in the case of individual service providers. The most favoured transition technique, known as “dual-stack” involves the parallel operation of both IPv4 and IPv6 in all components, applications, and services on a given network, for an extended period of time. This is in order to provide direct IPv4-IPv4 and IPv6-IPv6 connectivity wherever needed, without the use of inefficient translation techniques. By definition, dual-stack involves the

ongoing operation of IPv4, and therefore implies an ongoing requirement for IPv4 address space, for use in all dual-stacked infrastructure and devices. This is why even under the best circumstances, the extended lifetime of IPv4, whether through NAT or transfer mechanisms, is in fact a necessity for a smooth IPv6 transition.

However with an effective and wide-scale deployment of IPv6, the reliance on IPv4 in a dual-stack environment will steadily

diminish, since IPv6 is favoured wherever there is a choice. Across the Internet, IPv4 will diminish from being the dominant majority protocol, to being a secondary protocol (after IPv6), to being marginal, and then forgotten and entirely unused. The timescale of such a progression may be in the order of five to ten years from start to finish, by most estimates, assuming that IPv6 transition does continue steadily and at a sufficient pace.

---

## IPv6 and the Future Internet

It must be remembered that regardless of the advent of IPv4 transfers, and the possible emergence of an IPv4 market, the long-term future growth and success of the Internet is dependent on the successful deployment of IPv6. It is widely recognized that IPv6 provides the only means to achieve long-term growth while maintaining the critical technical features of the Internet: accessibility, neutrality and openness.

A failure of IPv6 deployment would not prevent global networks from continuing to grow, but would result in inevitable damage to or loss of these critical technical features. In particular, the widespread adoption of NAT, and CGN in particular, will steadily degrade the ability of any Internet device to connect to any other, and with it the global connectedness of the Internet. In turn, the applications used by devices will rely increasingly on application-specific

servers, which become less distinguished from the network infrastructure itself, compromising the neutrality of the Internet. Finally the barriers to entry for new service providers, including ISPs, and applications and content providers, will be raised by limited access to public IP address space, by the cost of effective NAT technologies, by the complexity of distributing application servers within the infrastructure, and by the inability to distribute content outside of these established applications.

*In such a future, the global network environment may still be referred to as “the Internet” but it would have lost the characteristics that have been critical to its success. This scenario is clearly not one for which we should plan.*

# Part 2. Background Reading

## What is IP?

IP is the Internet Protocol, a standardized set of communications mechanisms developed over the past 40 years. IP exists in the form of computer code in any device, service, or system that needs to connect to the Internet, and any implementation of IP needs to comply with its standards in order to work correctly and reliably.

As the name implies, IP is central to the Internet; in fact, its specific characteristics are central to the success of the Internet today. In particular, IP directly enables at least the following three essential features of the Internet:

- **Globality:** The ability to exchange data between any pair of devices at any pair of distinct locations (IP addresses) on the Internet
- **Neutrality:** The separation of applications and content from the infrastructure (the “IP layer”) of the Internet
- **Openness:** The minimal barriers to entry for technology, content, and service providers, who may freely access the IP standard specifications.

---

## IPv4

Version 4 of the Internet Protocol has been in use since 1 January 1983<sup>4</sup>. In the 30 years since then, IPv4 has been used to build the Internet as we know it, and the characteristics of IPv4 have defined the Internet in many ways, including the limits to its growth.

The addressing mechanism used by IPv4 involves a fixed-length 32-bit (binary digit) network address format, which provides a maximum of  $2^{32}$  (around 4 billion) unique addresses. This address format dictates

an absolute limit to the number of devices that can be directly connected to the Internet and is inseparable from IPv4 itself. When originally designed, this architecture seemed sufficient for all foreseeable uses; however, on today’s Internet, the IPv4 address space is no longer large enough. This has driven the adoption of its successor, IPv6, discussed below.

---

4. <http://www.iana.org/numbers>

## How “Big” is IPv4?

The limited size of the IPv4 address space is imposed by the 32-bit address format in the same way that an 8-digit format imposes a strict maximum of 100,000,000 unique telephone numbers.

While this strict architectural limitation may appear short-sighted, it is important to remember that many other networking protocols were under development at that time, including the Government-endorsed “OSI” system, and it would have seemed likely at the time that IPv4 would be replaced before its address space was exhausted.

Often the size of the IPv4 space is reported as 4 billion ( $2^{32}$ ) addresses, but because around 14% is designated for special purposes within the protocol (including “private addressing”, explained below), there are around 3.7 billion unique addresses available for use on the Internet.

For a variety of reasons, the utilization of IPv4 address space can never reach 100%; in practice, 50% may be more realistic. With an Internet user (and device) population now estimated as well over 2 billion, it should be clear why the days of IPv4 are known to be numbered.

## A History of IPv4 Management

The Internet requires IP addresses throughout the entire network, and these addresses must be allocated uniquely to each and every device that is connected. In addition, addresses must be allocated “hierarchically”, that is, in blocks that correspond to the Internet’s topology of networks<sup>5</sup>, and within those blocks to smaller sub-networks and infrastructure. In order to manage and track the entire address distribution process in a systematic way, the Internet has always required a “registry” mechanism in some form.

During its early years (the mid-1980s) the Internet was considered an academic and experimental activity, and not likely to be useful in the long term. During the 1980s there was a single IP address registry (the “Central Registry”), located in Los Angeles, USA; which evolved under a series of names including the “NIC” (Network Information Centre), the InterNIC, and finally the IANA (Internet Assigned Numbers Authority<sup>6</sup>). However within a few short years the Internet had grown beyond expectation, and many of its mechanisms, including the address registry system, needed to adapt to that growth.

By the early 1990s it was clear that a new approach to this central registry arrangement was needed. The Internet standards body, then IETF (Internet Engineering Task Force<sup>7</sup>) decided, through its “RFC” (Request for Comments) process, that the registry function should be

5. <http://www.ietf.org/rfc/rfc1338.txt>

6. <http://www.iana.org>

7. <http://www.ietf.org>



regionalized to allow both a distribution of the administrative function, and also a more careful stewardship of the limited IPv4 address space.

During subsequent years, Regional Internet address Registries (RIRs)<sup>8</sup> were established in a number of locations, in the following order:

- RIPE NCC (1992), serving Europe and parts of Africa and Central Asia;
- APNIC (1993), serving the Asia Pacific region;
- ARIN (1997), a distinct registry organization assuming responsibility for the Americas and part of Africa;
- Lacnic (2001), assuming responsibility for Latin America and much of the Caribbean; and
- AFRINIC (2005), assuming responsibility for the African region.

As each RIR was established, it assumed responsibility for its designated region, either from the IANA, or else from an existing RIR that had been fulfilling that function. Since 1992, this evolving collection of organizations (IANA and the RIRs) has shared responsibility for management of IPv4 address space globally, via community-defined policy, procedures, and accountability mechanisms, which are well established and documented.

It is important to note that successful IP address management, in terms of both Internet growth and stability, requires more than a successful registration function. IP address space must be distributed in a hierarchical manner that corresponds to the physical topology of the Internet, composed as it is of many individual, independent IP networks. Each network operator, whether large or small, must demonstrate need for an allocation from the responsible RIR, under policies developed collectively by the addressing community<sup>9</sup>. Those who cannot demonstrate this need can always obtain needed addresses from their own Internet service provider, under this hierarchical distribution model.

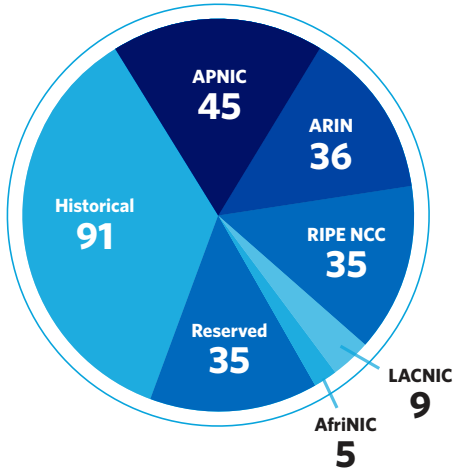
---

8. <http://www.nro.net/about-the-nro/regional-internet-registries>

9. <http://www.ietf.org/rfc/rfc2050.txt>

## IPv4 Address Distribution

In February 2011, when the last available blocks of IPv4 address space were finally distributed to the RIRs<sup>10</sup>, the global distribution of IPv4 address space was as follows:



Final distribution of IPv4 address space<sup>11</sup>, as at February 2011

It is notable that the “Central Registry” segment of the chart above accounts for some 35% of the entire IPv4 space. Known as “historical” or “legacy” address space, this was distributed prior to the RIR system, under generous and relatively informal allocation procedures that did not anticipate the Internet’s success. The rapid depletion of IPv4 address space by the Central Registry was one of the drivers for establishment of the RIR system<sup>12</sup>.

In the period since 1992 however, total Internet growth has accounted for over 99% of its current user population<sup>13</sup>, yet has been achieved using around 50% of the total address space. This indicates the success of the current distribution system, but also reveals that the Central Registry address space can be regarded as “under-utilized” by today’s standards.

## “Legacy” Address Space

As well as assuming responsibility for distribution and registration of IP addresses, each RIR also became responsible for the administration of legacy address space in each region. In most regions, early management policies may have distinguished between legacy and “current” addresses, for administrative and transitional reasons. However these policies have tended to converge, so that in most cases there is no longer any practical distinction between legacy and current space. This is an important issue in relation to IP address transfers, which are discussed below.

10. Internet Addressing in the 2010s

10. <http://www.nro.net/news/ipv4-free-pool-depleted>  
11. <http://www.nro.net>

12. <http://www.ietf.org/rfc/rfc1366.txt>

13. <http://www.ietf.org/rfc/rfc1366.txt>

## IPv4 Address Exhaustion

The final distribution of IPv4 address space from the IANA to the RIRs (in February 2011<sup>14</sup>) represented just one stage in the process of IPv4 exhaustion. This event left individual RIRs with IPv4 pools of different sizes, but did not impose any immediate change in their mechanisms of distribution from those pools. In accordance with their own policies however, each region has made a series of policy adjustments in response to the circumstances (as determined by community policy development processes).

In APNIC's case for instance, its remaining regular stock of IPv4 addresses was exhausted in April 2011, when it reached its "final /8" block of IPv4 (equalling some 16 million addresses). At this point, a pre-determined policy came into effect to impose an agreed rationing system, allowing each customer organization to receive only a single "/22" block (1024 addresses) from the remaining address pool. This approach, similar to that agreed by other RIRs, provides a small supply of IPv4 to new entrants from a supply that is expected to last beyond the time required for IPv6 transition (at least ten years).

Across the Internet, the exhaustion of IPv4 involves a set of ongoing processes, rather than a single global or regional event. Indeed, there was never any expectation or possibility of another outcome, due to the entirely decentralized nature of IP address management and utilization.

## Registries, Routing, and IPv4 Transfers

The role of an IP address registry is to record the authorized holder of a given address block; that is, the party who is responsible for its use. However the registry itself cannot strictly prevent any party from making use of a given block, with or without the permission of the address holder.

In practice however, an address block can only be used by any party on the Internet with the permission of others; specifically, an ISP that is asked to provide connectivity to the party, by "routing" their address block onto the Internet. An ISP that is asked to route a given block of addresses is obliged to identify the requestor and provide service only if their identity agrees with the registry record. That said, there are scenarios in which this may not happen, as follows:

- (1) "Squatting" refers to the unauthorized use of an address block, which is unallocated and therefore not registered in any registry.
- (2) "Hijacking" refers to the unauthorized use of an address block, which is allocated to another party. In case the address block is in active use, successful hijacking will cause service and connectivity disruption to the other (presumably legitimate) user of the address space.

---

14. <http://www.nro.net/news/ipv4-free-pool-depleted>

(3) “Black market transfer” occurs when a private agreement is made between the registered holder and another party, for an address block to be used by the party on a temporary or permanent basis.

An alternative to all three scenarios listed here is that of the “registered transfer”, which refers to a legitimate process of altering registration records through normal registry procedures. This is a part of normal registry activity, in cases of organizational mergers or acquisitions, and other administrative changes. Until recently however, registry policies have forbidden other forms of address transfer, for instance as the result of private agreements independent of the registration process (such as agreements to buy or sell address space).

## History of IPv4 Transfers

Prior to the exhaustion of available IPv4 supply, there was little incentive for transfer of IPv4 address blocks between independent parties. Those who needed

address space were able to get it through well-defined RIR procedures, and those who held unused space were obliged to return it to the RIR (though this has been a rare occurrence).

Furthermore, RIR policies have prevented the free transfer of address blocks, by disallowing registration of blocks “received” in this way. Therefore while a “black market transfer” could occur without registry involvement, the recipient would face difficulty in having such address space routed, and the risk of having it “reclaimed” by the RIR.

On the other hand, any successful use of black market space would have detrimental effects in terms of network security and engineering, because address users cannot be easily identified through the address registry. If such use were to become widespread, the integrity and value of the IP address registries would be seriously diminished, which would in turn have negative implications for Internet stability and security.

## IPv6

The success of the Internet has been increasingly apparent throughout its history, and during the 1990s it became clear that the IPv4 address space would not be sufficient to support Internet growth indefinitely<sup>15</sup>. While processes for managing IPv4 address space were being improved, another effort was underway to develop the successor to IPv4, known as IPv6. This was driven primarily by the need to accommodate long-term Internet growth within a much larger IP address space, though certain other technical features were introduced within IPv6 at the same time.

Importantly, it is possible for the Internet to continue growing with IPv4, by using a range of techniques that can extend the lifetime of that address space, features which are already being deployed. However in doing so, the Internet will steadily and inevitably lose the characteristics of globality, neutrality, and openness which were described earlier. IPv6 provides the means to retain those essential features while the Internet continues to grow into the distant future.

## IPv6 Address Management

With the adoption of the IPv6 protocols, the role of managing the IPv6 address space was delegated to the RIRs, under the same regionalized management system. This decision reflected the reality that IPv6 addressing has the same architectural constraints as IPv4, requiring the same hierarchically based management approach<sup>16</sup>; and the fact that there is no characteristic of IPv6 which suggests a different approach.

### How Large is IPv6?

The size of IPv6 address space can be described and has been reported in a variety of ways, but in terms of individual addresses the address size of 128 bits provides  $2^{128}$  different values, equal to  $3.4 \times 10^{38}$ , or 340 trillion trillion trillion addresses in total.

In practical terms however, the addressable unit of IPv6 is the “subnetwork”, or “subnet”, which accommodates 64 bits of the IPv6 address field. The IPv6 Internet is able to address  $2^{64}$  ( $1.8 \times 10^{18}$ ) subnets, and each of these may comprise a few devices, or many. Even with just one device per subnet, and an overall utilization efficiency of 0.1%, a far-future Internet of 100 trillion devices (1,000 for each of 100 billion humans) would consume  $10^{17}$  subnets, just 5% of the entire IPv6 space.

---

16. <http://www.ietf.org/rfc/rfc1881.txt>

## Actions for Governments

As discussed in this paper, the transition to IPv6 must be a critical goal for all Internet stakeholders. In practice, it has proven to be a difficult goal to achieve due to lack of short-term incentives for investments, which are difficult to justify in the highly competitive Internet environment.

For these reasons, positive intervention from Governments may be justified, not only for the sake of national interests and competitiveness, but also for the health of the Internet itself. While it is not the aim of this paper to provide comprehensive guidance in this area, governments and regulators may consider the following measures, in order to assist in IPv6 deployment in their jurisdictions:

- (1) Procurement: All government purchasing and contracting should place IPv6 compliance requirements on ICT products and services.
- (2) Implementation: All government online and Internet-related ICT services should support access via IPv6.
- (3) Standards: Technical and functional requirements should reference specific standards (RFCs) in expressing Internet-related requirements.
- (4) Certification: Governments can establish certification mechanisms, or recognize existing certification bodies, in recognizing specific products as "IPv6 ready".
- (5) Incentives: Financial incentives may be considered in the form of tax relief, or other assistance provided to those who undertake concrete transition steps within a required time frame.
- (6) Capacity building: Training providers and consulting firms may be encouraged to provide suitable IPv6-related services to government and/or industry members.
- (7) Monitoring: Governments may seek direct or indirect data sources, which help to monitor progress of industry and other sectors towards IPv6 goals.
- (8) Leadership: Government may promote IPv6 activities through media, events, competitions, awards, and other mechanisms.

## **For further IPv6 information visit**

<http://www.apnic.net/community/ipv6-program>



<http://www.apnic.net>

**APNIC**

