IPV6 DEPLOYMENT, NOT ONLY A NETWORK PROBLEM

Omar Walid Llorente, Tomás P. de Miguel Moro, and
David Fernández Cambronero
Universidad Politécnica de Madrid, Spain
omar@dit.upm.es, tmiguel@dit.upm.es, dfernandez@dit.upm.es

Abstract The new IPv6 and the current IPv4 will coexist for many years. A wide range of techniques have been designed to make the coexistence possible and to provide an easy network transition. An organization involved in the IPv6 transition should address not only network issues but also final user applications. The management of the different services must also be revised. This is specially relevant in academic institutions where educational, production and research networks live together in the same environment. This document is intended to give the reader a more comprehensive and accurate picture about the IPv6 transition procedures that may be accomplished by the different University members to introduce IPv6 smoothly and successfully.

Keywords: IPv6 transition, IPv6 addressing plan, IPv4/IPv6 integrated management, IPv6 campus transition

1.1 Introduction

The current version of IP has not changed substantially since RFC 791 in 1981. To address the requirements of the New Generation Internet, the new IPv6 protocol has been designed. The new IPv6 and the actual IPv4 will coexist for many years. A wide range of techniques have been designed to make the coexistence possible and to provide an easy network transition. However, when an organization adopts IPv6, many issues should be taken into account. Although IPv6 is a network protocol, transition is not only a network issue.

IPv6 is not an update back-compatible with IPv4, but instead a complete new network protocol oriented to radically change (or upgrade) the way our computers and electronic devices collaborate one with each other. There are a lot of new features related to the new protocol, but the more relevant ones are the huge addressing capability and the more efficient hierarchical addressing and routing structure. IPv6 is able to address 2^{128}hosts and networks, what comes to be approximately like our current public addressing capacity elevated

to the fourth power: $[Internet(IPv6)] \simeq [Internet(IPv4)]^4$. This large address space has been created to allow multiple levels of subnetting and address allocation, from the Internet backbone to the individual subnets within an organization.

Thanks to this range of available unique public identifiers, virtually everybody will be able to use end-to-end communications with their electronic devices, avoiding the current use of proxies and network address translators (NATs). The main benefits of the end-to-end model are threefold: the privacy is greatly enhanced; the performance of the communication is augmented; and the eventual network problems can be handled by both ends allowing the upper protocols to control the new situation in a way that better fits the user application (i.e. changing the data rate or transferring the connection to a new network interface).

If the question is about how will an IPv6 router will manage the huge number of networks that will be deployed with the current processing capabilities of the hardware, the problem is solved by using a more efficient hierarchical route path distribution: in the new network protocol, the network prefix will be assigned not to the final user organization, but to the Internet Service Provider (ISP) it will use. The new problem that arises is that the organization (and the final user) will have a different network prefix from each one of the ISPs hired by it. Network administrators will have to consider what solution fits their network best, among the many proposed.

1.1.1 Transition to IPv6

One of the main concerns about the new protocol is the way in which it can be introduced in existing organizations. IPv6 was thought to allow gradual migration, but protocol transitions are not easy, and the one from IPv4 to IPv6 is not the exception. Protocol transitions are typically deployed by installing and configuring the new protocol on all nodes within the network and verifying that all host and router operations work successfully. Although this might be easily managed in a small organization, the challenge of making a rapid protocol migration in medium and large organizations is very difficult. This is specially relevant in universities where advanced network researches are combined with administrative tasks and basic academic nodes, together in the same IP network.

It is then clear that the transition from the current IPv4 networks to the new IPv6 scenario will be a time consuming process during which both protocol versions will coexist. To put the new protocol in reality, there are many methods to choose from, either of which allows the access between IPv4 and IPv6 systems. The most recommendable ones are:

- IPv6 Protocol Translation (NAT-PT). Like IPv4 Network Address Translation (NAT), but with IP protocol translation also.
- Tunneling connections. This method allows the use of IPv4 packets to transport IPv6 ones and vice versa.
- Dual Stack Transition Mechanism (DSTM). Permits the use of both protocols at the same time by means of installing two different protocol stacks over the physical network drivers. The device is assigned an address of each type for each stack.
- Application Layer Gateways (ALGs). For example, an HTTP cache server that can fetch pages over IPv6 and transmit them to the client application using IPv4.

NAT-PT has at least the same constraints of the NAT approach in IPv4 networks, including requiring different translators for the different applications to be translated (i.e., one for HTTP, another for FTP and other for IRC) in the worst case. Although tunneling allows the user to have a real experience in IPv6 without noticing the intermediaries, it has various constraints from the administration point of view (like establishing and maintaining the tunnels, tracing the problems and managing the routing tables for the different connections, besides of the processing and memory requirements for the routers).

The ALGs, apart of interfering with the Application Level, may not allow the user a real IPv6 experience. On the contrary, DSTM is a method that does allow the user to feel a real IPv6 experience because it brings the new protocol to the user desktop with all the implied consequences. The main problem in DSTM is that it needs, for the Layer 3, two pieces of software capable of adequately managing both the old and the new protocol at the same time. This is its main advantage too: the user can still use both protocols simultaneously, each one with the required or supported applications.

1.1.2 Planning is the key

It is natural that the IPv6 transition process will be long in time. It is natural too that it won't be a costless one. Much less if the necessities of the organization do not fit within the initial planning.

As IPv6 transition planners, our first effort has to be oriented to being able to bring the new protocol to the users in an 'easy way'. Not only for them, but also for the network managers. For example, allowing the user to experiment the IPv6 benefits while using the same equipment or by reducing the number of changes that the internal and external networks will suffer.

As previously presented, there are a number of technical approaches for making the IPv6 transition available, each one with its own advantages and

drawbacks. But even if only one transition mechanism is chosen for all the network, many configuration alternatives are available. Besides, the actual network technology used, the services offered and the operating priorities in each Campus to access the Internet may introduce some other considerations in the sense of available bandwidth, traffic interference, operative system and application software upgrading requirements, functionality, etc.

All these points will draw a different roadmap on each case to achieve the aim of being IPv6 aware, but from our experience, this task will not be overwhelming if the initial planning criteria are good enough. We consider that IPv6 is something that the end users cannot skip, empowering not only them and their applications but their work too. Therefore, putting IPv6 on their hands has to be a clear priority and has to be carefully planned.

1.2 Network level transition

Although there have been a tremendous number of different physical network technologies used in Campuses all around the world, nowadays, from a physical point of view, and as a de facto standard, the Ethernet (802.3) and its derived or related protocols (802.11b, 802.11g, 802.1q) are established as the actual LAN and WAN protocols for the majority of the Universities and High Schools.

This consideration, in addition to the fact that the Dual-Stack option is widely available on most modern OSs (like Windows XP, GNU/Linux, FreeBSD and many others), if taken into account, can polarize our design in two different ways: 1) to add IPv6 to the currently deployed networks creating a full DSTM network, and 2) to create a new, isolated and initially experimental network tree where the IPv6-only option is considered[1].

Which one of these methods is the best for each case is a common sense decision that has to be made taking into account the special circumstances of each organization and its users. Our experience points out that both approaches can be used in the same environment if it is necessary, but complexity is a factor to decrease.

From our point of view, the best approach for the IPv6 Transition Process in academic organizations is the full DSTM one. The main reason is that with this method, it is always possible to deploy the new protocol step by step and network by network, and then slowly adapt each host to the use of both protocols by upgrading or configuring its Operative System. This offers more control on the situation to the administrative personnel and greatly reduces the initial costs of the IPv6 deployment.

[1] The 2nd option can be configured in several ways: 2.a) By using a different interface in each host to access the separate IPv6 network, and 2.b) by using the same interface but another logical LAN (if 802.1q or LAN tagging it is available in the network core) to access the separate IPv6 network.

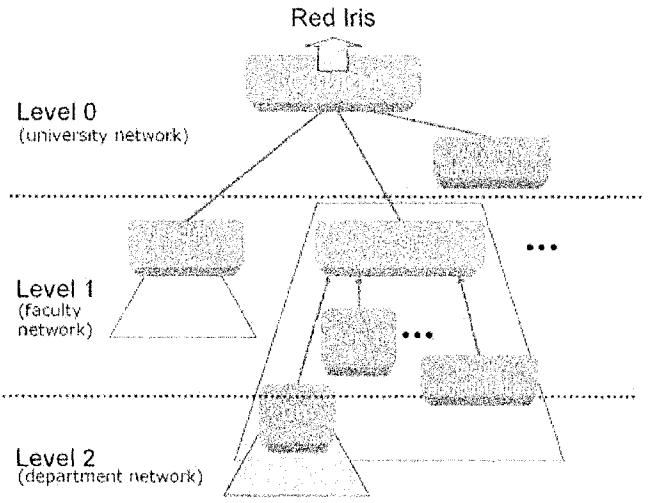

Figure 1. UPM network topology

1.2.1 Addressing plan

The Technical University of Madrid (UPM) network is organized as a star-shaped structure. Interconnection is provided by means of Giga-Ethernet links between all University High Schools, University Institutes and the Central Network Administration Department (CCUPM). Internet connectivity is provided by the Spanish Academic Network ISP (called RedIRIS) using another Giga-Ethernet link between CCUPM and the regional (Madrid) node of RedIRIS.

The UPM network is a multi-level topology (see Figure 1). At the first level is the CCUPM router which links University subnets with the rest of the Internet Academic Network. The second topology level is composed of all Faculty routers. Behind the Faculty level there is the department or research group topology level. This lower level is composed by one or more subnets connected through a firewall to protect internal users from external attacks.

RFC 3513 defines the IPv6 protocol addressing architecture. As it is described in section 2.5, an IPv6 address is composed of 128 bits divided into:

- a Subnet Prefix, that identifies the subnet (link, in IPv6 terminology) where the system is connected, and is common to all the stations connected to it.

- an Interface Identifier, which identifies a station inside a network, and must be different for every system connected to it.

RFC 3513 in section 2.6, states that, with very few exceptions, the interface identifier must be 64 bits long, and must be derived from MAC address using the EUI-64 format.

The large size of the interface identifier is not due to the number of stations connected to a subnet, as it happened in IPv4. The cause is the simplification of the autoconfiguration mechanisms derived from this length, since the generation of a unique identifier derived from MAC address (48 bits long) is very simple. Therefore, the autoconfiguration mechanism simplicty prevails against the efficient usage of bits, which is not critical in IPv6 due to the large addressing capability.

The RIPE-267 document defines the actual address assignment policy in Europe. Basically, it is proposed to assign 32 bits long prefixes (/32) to IPv6 Internet providers. Regarding prefix assignment to ISP users, the RIPE document follows IAB/IESG (RFC 3177) recommendations, consisting in:

- As a general practice, assigning 48 bits (/48 prefix) to each ISP user, except large size users who will occasionally receive shorter prefixes (/47 or even shorter).
- Assigning 64 bits (/64 prefix) only if it is justified by a design decision of assigning a single subnet, for instance, a mobile network inside a car or an ADSL residential network.
- Assigning 128 bit addresses (/128 prefix) in case of being sure of having a single device connected (for instance, a PPP connection through telephone network).

It is preferred to sacrify address space in order to improve the network management efficiency and to simplify the address management, which is very expensive in actual IPv4 based networks. The idea of having a single address allocation in most of the cases underlies this assignment policy, which would cover current and long term needs, against usual high cost reassignments performed nowadays.

In addition to that, the studies performed (RFC 3177, section 4) suggest that despite the high misuse of addressing space no addresses shortage will occur, even according to the more pessimistic forecasts.

According to the previous recommendations, RedIRIS has been delegated by RIPE 2001:720::/32 prefix. And, in the same way, RedIRIS IPv6 addressing plan establishes the assignment of a /48 prefix to each University. As the prefix assigned to UPM by RedIRIS is not known at the moment, it will be referred to in this document in the following way: 2001:720:XXXX::/48.

Once such network prefix is assigned to the Campus, the main objective has to be to assign subnetwork prefixes for the different Campus users in a way that could allow subsequent assignments to be contiguous to the previous

ones. This is important because managing one IPv6 network prefix is much simpler than administering and configuring more than one.

The way this is done is the following: to assign sparse prefixes to the subsequent hierarchical levels by leaving blank (unassigned) spaces between them. Of course, the network mask lengths have to be proportional to the number of final subnetworks and users.

If UPM is assigned a /48 prefix, there are 16 bits for subnetting that are available for making /64 final user subnetworks. These 65536 /64 subnets can be distributed in the following way:

- 64 very big schools if 6 bits for each school /54 prefix are used, having 4*256 possible /64 subnets into each school, or
- 128 big schools using a /55 prefix, having 2*256 possible /64 subnets into each school, or
- 256 medium size schools using a /56 prefix, having 256 possible /64 subnets into each school.

Studying the different school needs and according to our estimations, the best approach in our case may be the following:

- 16 type A schools, each one with a /53 prefix network (each center with 4 blocks of 256 /64 subnets assigned and 4 more blocks to assign in the future), and
- 16 type B schools, each one with a /54 prefix network (each center with 2 blocks of 256 /64 subnets assigned and 2 more blocks for future use), and
- 16 type C schools, each one with a /55 prefix network (each center with 1 block of 256 /64 subnets assigned and 1 more block not assigned), and
- 32 type D schools, each one with a /56 prefix network (each center with 1 block of 256 /64 subnets assigned).

Using this addressing structure, which represents a total of 48 + 32 different schools or centers (see Table 1), it is possible to give to the more than 32 UPM institutions an adequate IPv6 prefix and thus to all their current and future users, hosts and networks.

1.2.2 Routing issues

There are many IP routers that are IPv6 capable. Probably most of the ones deployed for IPv4 are capable of managing IPv6 packets by upgrading their software. And many times even without upgrading them. As a matter of fact, it

Table 1. UPM prefix delegation schema

Center type	Address format	Prefix Length
A	0CCCCBBBnnnnnnnn	/53
B	10CCCCBBnnnnnnnn	/54
C	110CCCCBnnnnnnnn	/55
D	111CCCCCnnnnnnnn	/56

may be quite inexpensive today (forgetting the administration costs) to deploy a PC based IPv6 router if an open-source OS -like FreeBSD, GNU/Linux or any other- is used for routing one or both protocols into our networks. Other commercial solutions of the most important firms are widely available too.

As previously stated, if Dual-Stack infrastructures are to be deployed, it may be crucial to have routers capable of managing IPv4 and IPv6 packets without loosing performance. If this is not achieved, our users may be obliged by the circumstances to avoid using IPv6, thus wasting all the efforts done.

The main activities that an IPv6 router has to perform, besides the usual ones like forwarding IP packets and re-distributing routes (Inter-domain Gateway Protocols, EGPs, and Intra-domain Gateway Protocols, IGPs), are to generate Router Advertisement (RA) packets. This is done in order to inform the hanging nodes which prefix is assigned to the network or subnetwork and to permit the Stateless Autoconfiguration of the IPv6 devices connected to take place.

Managing RA configuration is an easy task if only a router per network is deployed. In other cases, as in multihoming, it is necessary to configure the routers to export the correct prefixes and the nodes to import, understand and manage all the RAs learned.

In large organizations, the route distribution protocols (EGPs and IGPs) play a very important role. Usually, the EGPs are imposed by our ISPs. Therefore, few changes can be made regarding this aspect. What has to be planned in any case for the migration is the IGP protocol, if needed (many times, as is our particular case, the static-route deployment approach may be good enough if the organization is not too dynamic or unstable).

If the IGP is necessary, and the one used to manage IPv4 routes does not support IPv6, a decision has to be made: to deploy a different and independent routing protocol for IPv6 or to upgrade to a new IPv4 and IPv6-capable routing protocol. It is widely recommended, because of its mature state and high compatibility, to use IS-IS [IS-IS] as IGP protocol even when OSI environments and TCP over IPv4/IPv6 have to live together.

Depending on the particular details of the migration, it might be considered to have (even when deploying DSTM) two different routers, one for each IP

protocol. This approach permits to avoid collisions and interference between the two IP version packets, specially if the new technology is not well known or stable.

1.3 Services transition

Ordinary users see the Internet through the applications they use daily for their work, from electronic mail to the WWW navigation. In general, users get great benefits from all Internet services, even the simplest ones such as FTP or Telnet. This section deals with the concerns of services migration.

1.3.1 Transition is not only a network problem

Existing applications cannot use IPv6 without previous modifications. One of the reasons is that TCP/IP network architecture is not perfectly layered. For example, upper layer applications usually use IP addressing to identify destination nodes and their traffic flows. Although many times symbolic names are provided, standard communication libraries based on socket interfaces use only IP addresses. Therefore, IP address management should be maintained as part of the application and complemented with the Domain Name System (DNS) management.

Besides, the new IPv6 addresses' structure forces the use of a new transport layer interface (the IPv6 sockets layer interface). Hence, from the applications point of view, the IPv6 deployment requires changes in the existing code and -maybe- the addition of new communication design concepts. The necessary changes to allow the IPv6 operation of any networked application are not deep nor specially difficult but have to be done by the main developer team, or at least, if the source code is available, by any interested group in the matter (usually public founded projects, research teams or universities).

In the case of the usual academic application set, there are a lot of applications already ported or at least IPv6-capable. A very good link related to IPv6 applications and protocol support is [Bieringer, 2005]. For the reader to have a sight of the status, [IPv6 status] is recommended.

As a very short summary, the most important networked applications and protocols of our environment that lack IPv6 support are the following:

- SAMBA/NetBIOS[2]: File transfer protocol used for sharing data, authenticating and integrating Windows and Unix users, printers and computers.

[2]There is a non-official IPv6 patch available for 2.2.5 version of samba system, but it doesn't provide IPv6 interaction between GNU/Linux and Windows OSs.

- NFS: Networked filesystem used to share files efficiently in (originally Unix-like) local networks.
- Active Directory (AD): Windows based authentication and resource integration system.
- Vulnerability scanners: the most usual way of testing the exposure of our organization to security problems.

Other applications that do not have good enough support on IPv6:

- IP filtering: In a potentially aggressive environment such as Internet, some kind of adequate filtering policy has to be provided. But, in IPv6, only one stateful packet-inspection implementation is available for the task. Other solutions are still in an experimental stage. Options like translation at network level (NATs) or proxys at application level go against the end-to-end paradigm of IPv6 and should be deployed only in very specific environments.
- Event logging apps: Only one server (mysyslog) of the many Unix syslog protocol implementations has native IPv6 support. Maybe this is not a problem in a Dual-Stack transition approach.
- Webmail interfaces: In [IPv6 status] only one webmail implementation is IPv6 aware. This may be a problem if a fully functional IPv6-only network is the transition method chosen.
- Instant messaging (IM): currently only the IRC protocol is fully supported, but it is expectable that other popular instant messaging systems will be available soon. Euro6IX project ([euro6ix]) has this item among its many aims.

1.3.2 IP filtering, security and other issues

As previously stated, unlike with routing protocols, the IPv4 and IPv6 filtering solutions are nor widely available nor fully-operational in most of the cases. Nevertheless, security of the new network has to be at least as important as in the old one. Indeed if a mixed (dual-stack) environment is going to be deployed.

Stateful IPv6 packet-inspection filtering is nowadays only performed by Cisco products in the commercial world. In the open-source world this is done by the connection-tracking capabilities of netfilter ip6tables and USAGI kernel extensions, that point to be in experimental status. See [Bieringer, 2004] for a recent status report and detailed explanations.

Other security consideration that has to be taken into account is that, at the time of this writing, there are not any security auditing tools (like there are for

IPv4, as saint, nessus, etc) capable of testing IPv6 systems further, in addition to checking for open ports (like nmap).

Besides, if we consider that, because of the universal availability of global IPv6 addresses and their derivation from the -unique- physical interface MAC field and because the addresses would be neither masked nor shielded anyhow, it is possible to track the interface -and thus the user- all over the world since the day the interface is connected to the Internet for the first time.

To solve this problem, it has been proposed that some kind of random address assignation has to be provided. Some different approaches have already been proposed to do that IPv6 address assignation, but yet neither one is standardized or mandatory.

Of course, all these pointed problems and many more will be hopefully addressed in the future, but today they have to be subject of attention from the administration point of view and will be part of the job to implement the practical solutions for them.

Since the support of IPSec is mandatory in the IPv6 stack, a native virtual private network has been provided. The problem here comes from the point that if we agree that firewalls are needed for IPv6 and given that the end-to-end connection and the end-to-end security are mandatory, what kind of policy has to be deployed in order to permit or disallow encrypted connections that are opaque (IPSec tunnel mode) for the firewall? Here lies a big dilemma between allowing end-to-end encrypted connections from and to any IPv6 devices, or only to some ones that are previously known and allowed.

In other way, given the hierarchical routing structure designed for IPv6 and the fact that the network prefix will be no more owned by the final user nor by the organization but by its Internet Service Provider, it has been noticed that any networked organization that wants to have high network reachability will have to be connected to 2 or more ISPs, thus managing more than one global network prefix. The issues that cover this matter are commonly referred to as multihoming. Multihoming is not a new concept, but in IPv6 it has some implications that it does not have in IPv4 environments (derived from the fact that the prefix has to vary if the ISP changes). There are some experimental solutions proposed by different R&D teams that the initial planning of the IPv6 transition should study.

Support to mobility is one of the main advantages of the IPv6 protocol and it has a lot of implications in hardware to be deployed, inherited prefixes, published routes, end-to-end security and authentication, IP connection tracking and application independence of the real IP address used by the IP stack. Careful analysis is necessary if it is mandatory to have mobility in the new network.

Other services like Authorization, Multicast, Quality of Service (QoS) and Traffic Engineering (TE) techniques have to be adapted in many cases if the

network protocol changes, so a careful study of the available options has to be done.

1.3.3 install.hosts tool: the integrated management of IPv4 and IPv6 networks

It is important to notice that only upgrading or porting the applications will not suffice for supporting the whole new protocol. Management tasks will be duplicated if special efforts are not planned to minimize the administration of DNS records, along with firewalling policies and wireless infrastructure, among other duties. Our approach to this problem was to update our management tools to support both protocols and to configure the different services mentioned. Besides, the task of maintaining both networks (IPv4 and IPv6 based) can be overwhelming if any kind of automatic configuration is not used to simplify the human part of the job and minimize the human mistakes.

We have integrated the configuration of the IP assignation to a name in DNS with the management of DHCP/DHCPv6 systems and with the actualization of wireless infrastructure hardware (MAC) access lists, among other functionalities. This approach has been completely developed at the Department of Telematic Systems Engineering at the Technical University of Madrid, using a one-pass, perl based script system to which IPv6 capabilities have been added [Latorre, 2004]. The main script name is `install.hosts`, and its main database is text based and refered to as `tabla.numeros`. Its main functions are:

- To create the IPv4 and IPv6 direct/reverse DNS address records for ISC [ISC] bind servers deployed. Other DNS registers, like MX, CNAME, etc are configured as well automatically.
- To create the DHCP/DHCPv6 config files for stateful address configuration for ISC [ISC] DHCP server and DHCPv6 sourceforge [DHCPv6] servers.
- To upgrade the Wi-Fi infrastructure access lists (ACLs) for MAC address filtering.
- To update the IPv4 and IPv6 filtering rules and host/network databases (using fwbuilder [fwbuilder] XML format for both IP protocols), cross-check MAC and IP addresses for each filtered host and auto-disabling obsolete/non-related rules.
- To distribute all the changes to the different systems affected and to restart the services when needed.
- To update the IPv4 NIS server configuration.

- Other minor tasks, like /etc/hosts, /etc/ethers and /etc/networks files generation or automatic accounting reports.

The main advantage of the system is its simplicity, allowing to do many complex management tasks in a easy, efficient and quick way.

For example, to do the most usual tasks (simple tasks like adding or deleting one host to or from an existing network) network managers have to edit only the main database file and then run the main script. When more complex tasks have to be done (like creating a new subnet or domain name) it is necessary to edit the main script, the headers and the library files on which the system relays, and then run the main script. It is fair to say that although this work is not so easy and have to be done by specially skilled staff, the management system helps a lot comparing with the manual configuration of the related individual services. Fortunately, that kind of tasks are much less usual.

The possibility of doing the usual changes in an easy way can be seen as an added advantage of this integrated management system: work can be divided into profiles depending on the skills of the personnel without compromising the reliability of the management process.

1.4 Conclusions

The IPv6 transition is a process which ends when all nodes within a network install and configure the new protocol. Although this might be easily managed in a small organization, the challenge of making a rapid protocol migration in medium and large organizations is much more difficult. This is specially relevant in Universities where advanced network researches are combined with administrative tasks and basic academic nodes.

We can recommend four steps in the transition of Universities and Campus networks to IPv6: 1st) To set up dual stack on backbone routers; 2nd) Rationally, to extend dual stack to the end-user networks and hosts and to finish the corporative applications migration in the dual environment; 3rd) To start removing IPv4 from end-users subnets providing at least a tunneling method to maintain communication with old IPv4-only nodes; 4th) Finally, all systems are IPv6 and IPv4 is removed from all routers.

Some clues that have been presented in this paper are:

- IPv6 transition is a slow process.
- Users want to maintain old services.
- Users don't care about the technology itself, only about the benefits it brings for them.
- Migration is not only a network transition problem, but a service upgrading one. There are two main kinds of services to be migrated in

the process: the services oriented to the final users, and the services to manage and configure the network.

- Our experience is positive and this can only get better: Welcome to the IPv6 world!

References

P. Bieringer, "IPv6 & Linux". 1st Global IPv6 Summit. August 2004. http://www.bieringer.de/pb/lectures/PB-IPv6-Brazil-2004.pdf

Peter Bieringer page on IPv6, http://www.bieringer.de/linux/IPv6/

J. Bound (ed.), "IPv6 Enterprise Network Scenarios". IETF Internet Draft. July 2004. http://www.ietf.org/internet-drafts/draft-ietf-v6ops-ent-scenarios-05.txt

Current Status of IPv6 Support for Networking Applications, http://www.deepspace6.net/docs/ipv6_status_page_apps.html

DeepSpace6, http://www.deepspace6.net/

DHCPv6 project at Sourceforge, http://sourceforge.net/projects/dhcpv6/

A. Durand, S. Roy, and J. Paugh, "Issues with Dual Stack IPv6 on by Default". IETF Internet Draft. July 2004. http://www.ietf.org/internet-drafts/draft-ietf-v6ops-v6onbydefault-03.txt

Euro6IX: European IPv6 Internet Exchanges Backbone. http://www.euro6ix.org

Firewall Builder project, http://www.fwbuilder.org

C. Huitema, R. Austein, S. Satapati, and R. van der Pol, "Unmanaged Networks IPv6 Transition Scenarios". IETF Request for Comments 3750. April 2004. http://www.ietf.org/rfc/rfc3750.txt

Intermediate System to Intermediate System (IS-IS) IGP protocol IETF charter, http://www.ietf.org/html.charters/isis-charter.html

Internet Systems Consortium, Inc. (ISC), http://www.isc.org/

Latorre Sebastián, D. "Gestión de cortafuegos en redes departamentales IPv6". Career Final Project. ETSI Telecomunicación. Universidad Politécnica de Madrid. July 2004.

M. Lind, V. Ksinant, S. Park, and A. Baudot, P. Savola, "Scenarios and Analysis for Introducing IPv6 into ISP Networks". IETF Request for Comments 4029. March 2005. http://www.ietf.org/rfc/rfc4029.txt

MIPL: Mobile IPv6 for Linux. http://mobile-ipv6.org/

P. Nikander, J. Kempf, E. Nordmark, "IPv6 Neighbor Discovery (ND) Trust Models and Threats". IETF Request for Comments 3756. May 2004. http://www.ietf.org/rfc/rfc3756.txt

E. Nordmark, and R. E. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers". IETF Internet Draft. March 2005. http://www.ietf.org/internet-drafts/draft-ietf-v6ops-mech-v2-07.txt

M-K. Shin, Y-G. Hong, J. Hagino, P. Savola and E. M. Castro, "Application Aspects of IPv6 Transition". March 2005. http://www.ietf.org/rfc/rfc4038.txt

J. Wiljakka (ed.), "Analysis on IPv6 Transition in 3GPP Networks". IETF Internet Draft. October 2004. http://www.ietf.org/internet-drafts/draft-ietf-v6ops-3gpp-analysis-11.txt