

IPv6 : The Solution for Future Universal Networks

Sathya Rao

Telcom AG, Sandrainstr. 17
3007 Bern, Switzerland

Abstract. The communication networks and services are changing rapidly. The conventional circuit and packet switched networks are being replaced by next generation networks, primarily based on Internet Protocol. The rapid growth of web based services has led to the explosive growth of the internet. However, the current internet protocol (IPv4), which is the backbone of transmission control protocol (TCP/IP) networking, is rapidly becoming obsolete, with the inherent problems related with limited address space, security and QoS features. The new protocol IPv6 has been developed to overcome all these problems and to provide solutions for the next generation networks. This paper addresses the features of IPv6 Protocol, the status of standardisation, and various activities around the world.

1 Introduction

Europe's leadership in Internet technology and provision of user access should be based on an offering with unlimited address space, quality and security, properties the current Internet does not cater for. Europe should foster a unique leadership strategy in promoting the next generation networks based on the new Internet Protocol version 6 (IPv6) protocol in order to promote pan-European E-commerce, offering customer protection and benefits in terms of security and quality as services converge to run over IP. Such a Euro-IPv6 network will place Europe into a position of strength in comparison to the US with respect to New Internet technology.

The deployment of IPv6 requires a good spread of diverse technologies and the support of national Internet Service Providers across the whole European community. Expertise in these new technologies, which overcome the limitations of the current IPv4-based Internet, cannot be found in just a couple of European countries. The skills required lie in the areas of seamless deployment of IPv6 into a large existing IPv4 base, and provision of quality of service and security at host and gateway/router level.

The Internet has doubled in size every year since 1988. There are over 44 million hosts on the Internet and an estimated 200 million users world wide. By 2006, the Internet is likely to exceed the size of the global telephony network, should IP telephony have not replaced the existing telephony network by then. Moreover, tens of millions of Internet-enabled appliances will have joined.

The Telecom industry (manufacturers and operators) need to build strategies to cater for the mobile information society, deploying brand new products and services such as wireless Internet devices, Internet cell phones and personal digital assistants which will emerge to become the new telecommunications tool of the next decade. The mobile information society will need to deploy for this purpose IPv6 as a robust Internet foundation.

This strategy will get the European leadership entrenched in every aspect of the European industry in view of the explosion of the E-business and E-entertainment in Europe. It is estimated that commerce on the network (E-commerce) will reach somewhere between 1.7 T\$ and 3.0T \$ by 2003. That is only three years from now (but a long time in Internet years). Secure E-business and European privacy should be advocated and implemented at the network layer not just at an application layer. The security feature is built-into the IPv6 protocol to solve the issue, which is one of the major weakness of the current internet protocol.

2 Applications with IPv6 at Driving Seat

E-commerce will be a new driving force for new economies, creating new business sectors and new jobs across Europe. This new economy needs a robust platform to guarantee its success. The E-business and E-shopper should be able to gain the necessary confidence that they are doing business in a safe environment, which is not the case today. European E-commerce could back-fire in the mid to long term without adequate customer protection. E-commerce will also create the need for more address space and this new need is a healthy sign of the growth of the Internet in Europe, but then this growth has to be supported by guaranteed IP address space which is not available today.

The Internet is proving to be one of the most powerful amplifiers of speech ever invented. It offers a global megaphone for voices that might otherwise be heard only feebly, if at all. It invites and facilitates multiple points of view and dialogue in ways not possible through traditional, one-way mass media.

The Internet can facilitate democratic practices in unexpected ways. The proxy voting for stock shareholders is now commonly supported on the Internet. Perhaps we can find additional ways in which to simplify and expand the voting franchise in other domains, including the political, as access to the Internet increases.

The Internet is becoming the repository of all we have accomplished as a society. It is becoming a kind of disorganized Boswell of the human spirit. Shared databases on the Internet are acting to accelerate the pace of research progress, thanks to online access to commonly accessible repositories.

3 Ongoing Activities

The US government is funding the 6REN/6TAP testing native IPv6 as well as Internet2, a project that tests the impact of QoS and higher bandwidth on the current Internet. Internet2 has subscribed now to IPv6 as the result of their tests that higher bandwidth is not the only solution but a smarter packet is needed to achieve a better quality of service.

The Japanese government is funding the Wide Project which is a copy of the 6REN/6TAP initiative to take Japan into leadership in the New Internet. The Japanese government is pushing for active IETF involvement in order to secure RFC adoption or early RFC influence.

The IPv6 Forum has been formed recently to promote wide adoption of IPv6 specifications in developing next generation network products and services. The IPv6 Internet Initiative is a key milestone for a range of products and services under definition within the mobile information society platform. European Concepts based on GPRS, UMTS and 3G products and services depend dramatically on the deployment of IPv6. The convergence is an opportunity for the European switch manufacturers to take leadership into the New Internet and define Core Switch/Routers.

Within the European Union 5th framework 'Information Society Technologies' framework, a project named 6INIT has been started to promote the deployment of IPv6 networks and services, in collaboration with Japanese and Canadian partners.

Eurescom has initiated a project (P702 : Internet Protocol Version 6 - new opportunities for the European PNOs) to investigate the usage of IPv6 networks to replace the conventional networks for delivering conventional and future public services. Eurescom has also initiated a new project (P1009) to study the deployment and transition strategies for services on top of IPv6, e.g. Mobility, QoS support, Multicast Implementation, Network configuration and management.

4 Standards Status

4.1 The Internet Engineering Task Force and IPnG

The IETF (Internet Engineering Task Force) is very active in promoting IPv6 standards, through their Request for Comments (RFC) documents which are generally adopted as standards for implementation. IETF has constituted a special group IPnG (IP next generation) to promote IPv6 activity.

The current version (4) of the Internet Protocol (IPv4) uses node addresses that are allocated from a 32-bit space. This 32 bit address space is further classified to provision Class A, B and C ranges, which constituted network part of 8, 16 or 24 bit, with corresponding host part of 24, 16 or 8 bit, depending on the number of expected hosts on a given network. This led to inefficient use of the 2^{32} possible addresses, since many 'important' organisations automatically asked for class A or B addresses using

up 2^{24} or 2^{16} addresses at each single assignment, even when they often only had several host computers, or had many subnets with several computers on each. A second problem was that addresses were rarely re-claimed after they were no longer in use.

The terms of reference for the working group is maintain all good features of current protocol specifications, and enhance the features to guarantee smooth transition to next generation networks. The IPv4 protocol is simple, binds multiple protocols, simple management, but limited with scalability in terms of address space, topological flexibility, Quality of service support, security, etc..

IPv6 is planned to support very high speed (Gbps), range of subnets, low information loss and will function independent of media (terrestrial, mobile, radio and satellite), provides auto configuration possibility, high security for business applications, application specific QoS, and multicast addressing facility. IPv6 will be also backward compatible to work with the current IPv4 protocol (through tunnelling mechanism).

4.2 IPv6 Features

The **next generation networks** based on IPv6 will provide:

- 128 bit wide address space to cover all possible appliances connectivity
- Differentiated Services in terms of quality (bandwidth guarantee and transit delays for real time flows).
- Security in terms of access point authentication, message integrity and privacy.
- Auto-configuration and reconfiguration capabilities allowing easy modification of network architectures.
- Management facilities allowing the setting up of on-demand services and providing ISPs with accounting capacities.
- Wide range of applications and services.
- Mobile host capabilities allowing provision of transparent access whatever the physical access used, supporting the evolving UMTS capabilities, will be the issue of co-operation between the mobile IP related projects (e.g. WINE).

4.3 Activities in the IETF

Within the IETF now, detailed work on IPv6 specification is pursued. Changes to routing protocols, transport protocols (the pseudo-header checksum in TCP and UDP) and applications that reference IP addresses (particularly DNS, but also FTP) have been specified. More subtle work in routing (beyond OSPF and RIP v6 changes) needs to be done, and more especially, the impact on RSVP and on multicast routing and Mobile IP routing as well as RTP/SDP and MPLS needs a lot of work to see what the real benefits may be.

The critical missing piece in IPv6 is a deployment plan that includes seamless interworking with IPv4, but provides clear benefits to a site to migrate. Three possible ways this may happen are: VPNs, Satellite IP (DBS) and large scale PDA/GSM mobile phone integration by a provider and vendor. The importance of the seamless interworking is because the Internet is now far too large to envisage even the switch-over that occurred in going from previous NCP to IP in 1980; and that switchover was even painful then - with a few hundred hosts rather than tens of millions. Moreover, there is clear reluctance by commercial vendors to invest into extensive upgrade; it must be made worthwhile by the quality of the benefits provided.

4.4 Available Implementations, Products, and Services

We can divide implementations into host and router side code. In the router side, most of the major vendors have at least beta products for the basic IPv6, although it's not clear if their routing protocols change they still interwork with the legacy IPv4 networks. On the host side, major operating systems such as Windows 2000 will have IPv6 in, although to date, only the research part of Microsoft have released a Windows implementation. In a recent IPv6 Summit Microsoft announced their commitments to IPv6 and officially released the IPv6 protocol stack. Similarly CISCO also announced their commitment to the IPv6 networking, which provided early boost to the next generation networks world. For Unix systems, there are releases for most major flavours, although many are very early code, and have a number of shortcomings. The best systems are the public domain offerings for FreeBSD and Linux, including DNS for IPv6 and other important infrastructural tools.

The implementations have many of the components that have been defined - but not all. For example, strong security is mandatory in IPv6; political considerations related to export controls have made it very difficult to have exportable implementations which meet the Standards. Moreover, some of the key components, like IP Multicast, Quality of Service facilities and Public Key Infrastructure are not yet fully standardised.

4.5 Address Space in Reality

In practice, the IPv4 address space has lasted longer than expected due to two technologies: Classless Inter-Domain Routing (CIDR), and Network Address Translation (NAT). This has reduced the urgency to move to IPv6.

CIDR is a generalisation of the class based address assignment that was originally devised for IPv4. Nowadays, the address assignment authorities (and they are devolved to regions of the Internet geographically now) assign IPv4 (and IPnG (or IPv6 as it is known to some)) addresses hierarchically, together with masks. The mask

determines how many bits of the address are network and how many are host, and the mask can be different at different places in the network topology - this allows the address + mask to be treated like a variable length prefix. The forwarding decision that used to be made by routers, based on simple best-next-hop, is no longer a simple lookup; it consists of a longest-match procedure. Routing protocols no longer exchange lists of network numbers to build upon a network map, but now exchange addresses + masks, to allow this hierarchical address space management to work. A secondary, but very important side effect of this is that the routing tables can now be summarised; typically, a country might be assigned a short mask, and within the country, each region, longer masks. This allows each router nearer to each local region to contain small number of (even just 1 per interface) entries.

The questions of the technical and political feasibility of address-space deployment are relatively separate. One advantage, for example, of the larger address space is that mobile users can keep their same low-order addresses, while the mobile network operator controls only the high-order bits. However when there is a real conflict between technical and commercial pressures, the solution is less clear. It would be possible to carve out a complete set of numbers and address for IP-telephony; early attempts to do this in conjunction with the ITU telephony groups have failed so far. One suspects that this is partly because attempts to make IP numbering as universal and well-structured as the telephone numbering is seen as a very real threat by the PNOs to their telephone revenue.

NATs are another technique for containing the growth of address use, but are based in two other important requirements: to avoid having to renumber hosts, and to provide some network security. Many sites had assigned IP addresses in isolation from the Internet, and had used addresses already in use in the world-wide Internet. To avoid the problem of re-numbering all their hosts and routers, such sites developed a technology to translate 'on-the-fly' the addresses of source hosts within IP packets being forwarded through firewall routers. This was done only for hosts which wished to communicate with the 'outside world', changing the addresses in packets to and from them, from the internal non-unique ones, to ones drawn (dynamically assigned) from a small pool of legitimate public addresses. This works well in typical large corporate networks as few hosts communicate with the outside world at any one time (principal of locality!). Often, the dynamic assignment in the firewall router was controlled by an application level authentication protocol (e.g., based on an RPC mechanism, or even just a telnet/login user-name + password to the firewall router). This meant that the NAT acts as a quite effective barrier to outside hosts accessing internal hosts (even if packets could get in by pure random luck, the responses would not get out...).

4.6 IPv6 Roadblocks

Many of the improvements promised for IPv6 have been specified in a simplified way for IPv4; the specifications are often less powerful than is possible with IPv6, but would provide enhanced services. It has turned out to be very difficult to organise even the IPv4 deployments of IPSEC, QoS and Mobile IP. This is partly from lack of motivation by the many smaller ISPs, but it is also because of the need to orchestrate the introduction of the services. There is often little advantage in introducing such services in an isolated part of the system. Moreover, in some cases the specifications are not really complete, or it is felt that their viability yet to be demonstrated by the R&D community. Here there is the serious problem that small-scale deployment is no feasibility demonstration; large-scale deployment is often beyond the means of the R&D community, and is not really supported by those providing the research networks.

It is probable that these improvements will come more with IPv6 than earlier, because the whole of the transition to IPv6 must be managed to some extent in any case. This will encourage the establishment of larger testbeds, at least by the bigger ISPs and PNOs. There are some substantial test-beds already; who are active on the 6Bone, which is a set of European sites who are experimenting with IPv6 in an encapsulated form. During the last summit NTT (from Japan) announced the world's first ISP to support IPv6.

5 Evolution Scenarios

The Internet engineering community is promoting a new version of the IPv6 as the answer to the address shortage predicted for the current Version 4. IPv6 offers enough addresses that every computer, cell phone and set-top box can be hooked up to the 'Net. However, migrating a large network to IPv6 is so difficult that few organizations have committed to it.

Theoretically, Version 4 could support up to 4.2 billion devices, but the allocation of those addresses has not been very efficient. An attempt has been made to increase the efficiency with interdomain routing and allocation rules that go along with it. But the side effect of those rules is the proliferation of network address translation [NAT] boxes, which take a single Internet address and multiplex it among a bunch of different devices. It's a fairly ugly process from an architectural point of view, although it turns out to be very effective, and a lot of people are relying on it. But because NAT intervenes at the IP address level, it has some consequences for end-to-end security and integrity of the traffic

The key transition objective is to allow IPv6 and IPv4 hosts to Interoperate. A second objective is to allow IPv6 hosts and routers to be deployed in the Internet in a highly diffuse and incremental fashion, with few interdependencies. A third objective is that the transition should be as easy as possible for end-users, system administrators, and network operators to understand and carry out.

Probably the most straightforward way to introduce IPv6-capable nodes is a dual stack approach, where IPv6 nodes also have a complete IPv4 implementation as well. Such a node, referred to as IPv6/IPv4 node in [RFC 1993], thus has the ability to send and receive both IPv4 and IPv6 packets. When interoperating with an IPv4 node, an IPv6/IPv4 node can use IPv4 packets; when interoperating with an IPv6 node, it can speak IPv6. IPv6/IPv4 nodes must have both IPv6 and IPv4 addresses. They must furthermore be able to determine whether another node is IPv6-capable or IPv4-only. This problem can be solved using the DNS, which can return an IPv6 address if the node name being resolved is IPv6 capable, or otherwise return an IPv4 address. Of course, if the node issuing the DNS request is only IPv4 capable, the DNS returns only an IPv4 address.

In the dual stack approach, if either the sender or the receiver is only IPv4-capable, IPv4 packets must be used. As a result, it is possible that two IPv6-capable nodes can end, in essence, sending IPv4 packets to each other. This is illustrated in Figure 1.

The target scenario is to have an IPv6 backbone network, which can also provide seamless interconnectivity with legacy IPv4 network. The typical network scenario with IPv6 backbone is shown in the Figure 2.

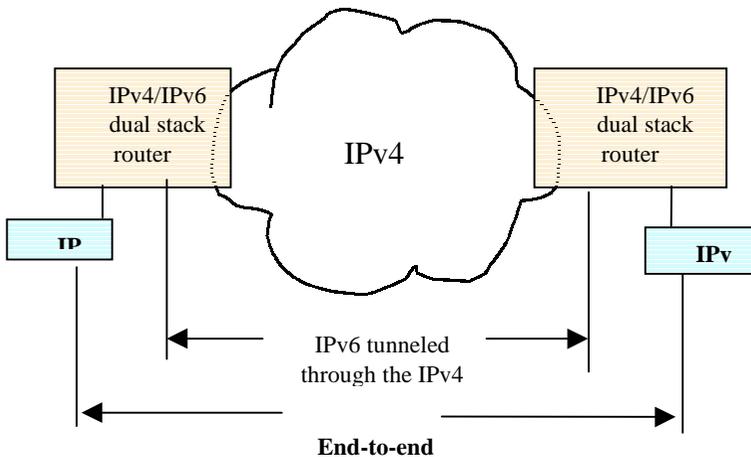


Figure1: A dual stack approach

6 Mobility and Internet

Europe is very strong in mobile networks deployment and usage. The internet access through mobile terminals and having an interactive data communication is a priority area in Europe. The first such applications already being introduced based on WAP.

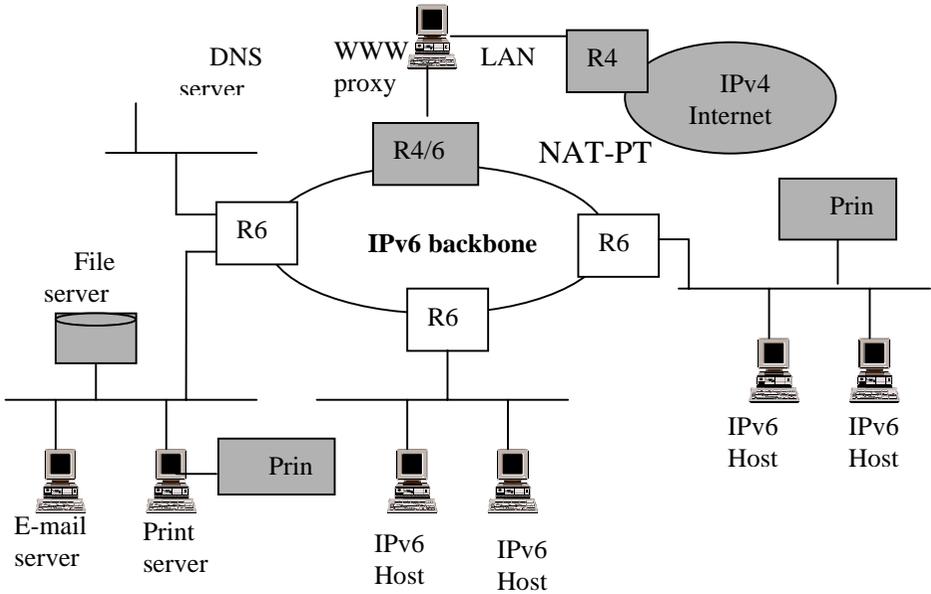


Figure 2: IPv4 and IPv6 network interworking

The WAP based communications are slow due to bit oriented non-efficient WAP protocol, though it provides the transition step for introducing mobile internet to the market. Third generation networks deployment plans to implement UMTS services are already in advanced stage of realisation. To progress the evolution and to enhance both mobility and internet features, the new initiative has been put in place called Third Generation Partnership Project (3GPP). The 3GPP is a global standardization initiative that was created just over a year ago, in December 1998, to produce technical specifications for Third Generation Mobile System based on the evolved GSM core networks and a new radio interface (UTRA). Major important steps have been achieved since then, such as the approval of Release '99 specifications in December 1999. The 3GPP work plan for the year 2000 includes Internet Protocol (IP) based communications. It is expected that as mobile phones gain access to Internet services, there will be an unprecedented growth in the demand of new Internet addresses as well as easier administration and tighter security. Convergence of Internet and Mobile Telecommunications move a step closer since IPv6 Forum has joined the 3GPP as a Market Representation Partner.

7 Conclusions

The specifications for a Next Generations Internet are largely complete from a technical viewpoint. There are still some loose ends, but they are not very significant. The large-scale deployment of many of the newer features over the current Internet is proving difficult, and will probably never happen on a large scale. Implementations of the basic feature sets are available in research prototypes, and starting to become available in commercial offerings. Implementations of advanced features are becoming available in research prototypes, but still require substantial experimentation and refinement. However, there are starting to be a number of substantial research networks on which the implementations are being deployed for R&D purposes. The general commitment to large-scale commercial deployment, and the time-scales over which this could be achieved, are still under discussion, though the recent announcements from the major vendors has brought the time scales to near short term plans.