

IPv6 standardisation issues

Jan-Pascal van Best
Delft University of Technology
PO Box 5015, 2600 GA Delft, The Netherlands
janb@tbn.tudelft.nl

Abstract

In the beginning of the 1990s, the IETF, the organisation that develops the Internet standards, started work on the successor of the Internet Protocol (IP). IP is the protocol that makes it possible for data packets to travel over the Internet from one attached computer to another. A decade later, the specifications for the new standard (named IP version 6, or IPv6) are almost, but still not quite, complete. The first commercial products and services based on IPv6 have just very recently been introduced into the market. In this paper possible reasons for this relatively slow process (compared to the reputation of the IETF) are discussed.

Contents

1	Introduction	2
2	IPv6 Timeline	2
3	Why the lengthy process?.....	3
3.1	Possible Causes.....	3
3.2	Sources of information	4
4	Conclusions	6
5	References	7
	Appendix: IPng Working Group meeting discussion topics	9

1 Introduction

The current Internet infrastructure is built on the underlying protocol named IPv4 (Internet Protocol version 4). IPv4 prescribes how data packets are sent over the Internet from one host (computer attached to the Internet) to another. IPv4 was standardised[1] in 1981 by the Internet Engineering Task Force (IETF, [3]) which sets the Internet standards. One of the things that IPv4 defines is the format of Internet addresses. In IPv4, Internet addresses consist of 32-bit numbers. This allows for a total of around 4 billion addresses. This limitation, further restricted by the way IP addresses are distributed over Internet Service Providers (ISPs) and the rapid growth of the number of hosts on the Internet, will cause the available IP addresses to run out. This is estimated to happen between 2005 and 2008. This address shortage was and still is the main reason for the development of a next-generation Internet Protocol.

The development of a successor to IPv4 started in 1990, but now, in 2001, it is still not finished. Internet Protocol Version 6, or IPv6, as the successor has been named, incorporates at least a greatly expanded address range. IPv6's 128-bit address space provides billions of addresses for every square meter of the planet Earth, so the main goal of IPv6 will at least be met. Other improvements will be in ease of routing, lowering the workload on the Internet's switchboards, authentication, privacy and quality of service. A problem hampering the introduction of IPv6 is that it is not interoperable with IPv4: you cannot use IPv6 to address a host that only knows IPv4 (although you can use IPv4 and IPv6 simultaneously on the same host). This means that any host on the Internet will need to use both IPv4 and IPv6 until all of the Internet has switched to IPv6. This, in turn, means that a host only needs to switch to IPv6 when all other hosts have switched, because all of them will be reachable using IPv4 until that time. A chicken-and-egg situation in exponential form. The standardisation process (up to now) of IPv6 is described in more detail in Section **Error! Reference source not found.**

Meanwhile, although not yet half of the 4 billion odd IP addresses have been allocated, the shortage of IP addresses is being felt in the Internet world. It is getting harder for ISPs to obtain IP addresses from the registries. In turn, the ISPs save on the use of IP addresses by using dynamic IP allocation (which gives their customers a different IP address on each connection) or even by using Network Address Translation (NAT). With NAT, customers do not get a real IP address at all, but they are given an address from the ISPs private address space, which is only connected to the Internet by a proxy or an address translating server. These ISP techniques contribute to asymmetry in the Internet, because they make it harder for customers of an ISP to, e.g., run their own web servers and many types of peer-to-peer internet applications.

Remains the question what has caused the lengthy process that should lead to a renewed Internet. Possible causes will be identified and discussed in Section 3, followed by a concluding Section 4.

2 IPv6 Timeline

In this section we will discuss the most important milestones in the IPv6 standardisation process.

1992: IETF issues Call for IPng proposals

The IETF started its effort to select a successor to IPv4 in late 1990 when projections indicated that the Internet address space would become an increasingly limited resource. This looming lack of IP address space was from the very beginning the most important driving force behind the development of IPv6. In July 1992, the IETF issued a Call for Proposals for a next-generation Internet Protocol, to solve the address space problem. Several (7+) proposals sprung up, with a wide range in revolutionariness. Late 1993, the IETF formed the IPng (IP next generation) Area (later transformed into the IPng Working Group[4]) to investigate the various proposals and recommend how to proceed. This IETF Area produced an IPng technical criteria document, RFC 1726[5] that was used in order to evaluate the three remaining proposals. The main criteria named in this document are:

In the late spring of 1992 the IAB issued "IP version 7" [...] recommending "an immediate plan for using CNLP as the basis for IPv7." After spirited discussion, the IETF decided to reject the IAB's recommendation and issue the call for proposals [...]. (RFC 1752)

- Architectural simplicity (don't include functions into IP that are better located elsewhere)
- Scale (allow for at least 10^9 leaf networks)
- Performance (routers must be able to process Internet traffic using this protocol)
- Transition (straightforward transition plan from IPv4)
- Configuration ease (provide for automatic configuration of hosts and routers)
- Security (provide a secure network layer)
- Service classes (associate packets with particular classes of service and treat them accordingly)

In RFC 1752[6] (January 1995), the then remaining IPng proposals were evaluated. One of the proposals (SIPP), with a number of modifications, including the doubling of the proposed address length from 64 to 128 bits, was chosen as a basis for IPng.

1995: RFC-1883 -- Internet Protocol, Version 6 (IPv6) specification

Based on SIPP, RFC 1883[7] was the first specification carrying the name IPv6. This specification has the IETF status of "Proposed Standard". Accompanying and following this specification, several other RFC's were issued, specifying the IPv6 addressing architecture, changes to related protocols, network autoconfiguration, security and IPv4-IPv6 transition mechanisms.

1998: RFC-2460 -- Internet Protocol, Version 6 (IPv6) specification

It took until December 1998 for the RFC-1883 based specifications to evolve into RFC 2460[8] (and accompanying RFCs), an IETF "Draft Standard", one step further on the way to the status of "Internet Standard". At the moment, the IPv6 specification is still incomplete: some parts that need to be specified in accompanying documents are not yet finished. There is still considerable discussion, especially about the IPv6 addressing architecture.

2001: First commercial IPv6 offerings

Only very recently Cisco Systems, the world's largest maker of data networking equipment, announced a range of software for its routers using IPv6[17]. It will be available at the end of May. Microsoft only expects full commercial support for IPv6 in its Windows operating system in 2002. Japan's NTT was the first service provider to offer commercial IPv6 services, undoubtedly pushed by the Japanese government, which pledged to adopt IPv6 by 2005. Zama Networks, based in Seattle, was the first U.S. service provider to offer commercial IPv6 services, starting operations in the spring of 2001.

In the following section, we will take a look at possible causes for this prolonged process. The IETF is supposed to be able to quickly find working solutions to engineering problems. Then why isn't this problems solved after more than a decade?

3 Why the lengthy process?

In this section, we will discuss a number of possible causes for the lengthy process of finding and deploying a successor to IPv4. Then we will search historical material for support for this possibilities, and discuss the results.

3.1 Possible Causes

In the field, several causes have been suggested for the long time scale of the IPv6 process. In this paper, we will consider three likely candidates and see whether these possible causes can be supported by historical material.

1. Proliferation of requirements

Some say¹ that the IPv6 standard has become too widespread and too complicated. This would presume that the expectations or requirements for the next-generation IP protocol have been extending during the standardisation process. Obviously, when extra requirements are constantly added during the process, the destination will never be reached.

2. No time pressure

Seeing that the first commercial products supporting the IPv6 standard have only come out recently, it seems reasonable to assume that there has been little pressure on the IPng working group to make haste. After all, the IPv4 address space is expected to suffice at least until 2005. Even though products are available now, vendors have seen almost no interest from their enterprise customers[18].

3. Academics and other “outsiders”

Has the openness of the IETF meetings and procedures given too much opportunity to relative outsiders, especially from academia, to interfere in the process, nit-pick on details and try to make the protocol of a universal quality, opposed to the usual IETF motto of “rough consensus and running code”? Have political motives, so despised by IETF regulars, seeped in with representatives from hardware vendors? In other words, have people from other backgrounds than usual at earlier IETF processes caused the retardation?

3.2 Sources of information

In order to check whether these three possible causes can be supported by empirical material, we have used a number of sources. First of all, RFC’s (documents issued by the IETF) regarding the IPv6 process have been studied. These include both RFC’s that define Internet standards (RFC’s of the *standards track* category) and other RFC’s (mostly of the *informational* category). Apart from these sources, interesting material can be found in the minutes of the IPng working groups meetings and in the archives of the IPng mailing list. Lastly, we have looked for articles in the technology press regarding IPv6.

RFC’s

First looking at the RFC’s, the main requirements for IPv6 have not changed between RFC 1883 and RFC 2460. These requirements correspond to those listed in RFC 1726. In order to see which of the accompanying RFC’s are considered to be part of the base IPv6 specification, together with RFC 1883 and 2460, we have looked at the references to external documents in the base IPv6 RFC’s. These references are listed in Table 1. This table gives an indication of what capabilities of IPv6 are considered as requirements.

	RFC 1883	RFC-2460
Format & Semantics of IPv6 addresses	X	X
ICMP for IPv6	X	X
Authentication extension header	X	X
Encapsulating security payload extension header	X	X
Security Architecture for the Internet Protocol	X	X
Use of Flow Label field	X	X
Path MTU Discovery	X	X
Jumbo payload extension header	(internal in RFC)	X
Use of Traffic Class field	-	X

Table 1: References to external documents from base IPv6 RFC’s

We see that there is little difference between RFC 1883 and RFC 2460 regarding their external references. The specification of the Jumbo payload extension header (related to very large IPv6 datagrams) is taken out of the base RFC and described in more details in a separate document. More importantly, in RFC 1883 the use of the Traffic Class field (there called Priority field) is described like it was in IPv4, in which it only labelled Internet traffic as non-interactive, interactive, or control traffic. In RFC 2460, the use of the same

¹ For instance Mike O’Dell, a long-time contributor and critic to the IPng effort

field (now named Traffic Class) is said to be under discussion and to be defined in other documents, with a reference to the “differentiated services” way of providing Quality of Service as a likely use for this field.

For the first possibility, this leaves us with little or no support. Even the use of the Traffic Class field is not discussed in the IPng working group, since the IPng working group does not see that as its role. And, as these RFC’s are technical documents, there is also no support for the second or third possibilities.

IPng Working Group meetings

In order to determine a possible change of focus of the work in the IPng working group, we have compared the topics in IPng working group meetings around the time of the issuing of RFC 1883 in 1995 ([9], [10]), of RFC 2460 in 1998 ([11],[12],[13]), and the two most recent meetings ([14],[15]). The results are in Table 2 in the Appendix. Explaining all the topics in the table would be impractical, so we will not do that here; most topics have lead to one or more RFC’s.

The first impression that Table 2 leads to is that many new discussion topics have been introduced between 1995 and 1998, but that from that time until 2001 only refinement and work on known problems has taken place. This would correspond to a remark in [13], where it is stated that “IPv6 is done”. Since most of the topics introduced after 1995 have in time lead to RFC’s, while those issue are not discussed in 1995, nor are they mentioned in the 1995 round of RFC’s, the conclusion seems justified that, in the period 1995-1998, the expectations of the results of the IPv6 working group have indeed extended considerably.

Regarding the second hypothesis, that lack of time pressure played a part, there is neither evidence for or against to be found in the meeting minutes. There are no remarks like “Since we have plenty of time, let’s...”, but also no-one says “We need to hurry, because...”. The third hypothesis finds also no support in these minutes.

IPng Working Group mailing list

As another source of information, the mailing list archives of the IPng working group have been studied. Special attention was given to three periods: the first month of the mailing list in August 1994; the months October-December 1995, around the time of issuing of RFC 1883; and the month December 1998, around the time of issuing of RFC 2460. The reason for examining the mailing lists around the time of issuing of major RFC’s is that it was expected that, shortly before such an RFC is issued, all the issues that should be handled in the specification are discussed, while in the period directly following a roadmap is set out of what still has to be done.

In the first month of the mailing list’s existence, most of the topics that are discussed in the working group meetings around December 1995 are already present. Noticeably missing during this month is discussion about mobile hosts and relations of IPng to application level protocols (such as FTP). More interestingly, a lot of heated emails discussed the addressing architecture, and especially the use of ISO-style addresses within IPng. ISO proponents would like the IPng addressing structure to be incorporated in ISO’s NSAP structure, or even directly use 160-bit NSAP addressing for IPng. Since there had already been a decision for 128-bit addresses, structured by Internet provider hierarchy, the IPng working group regulars seemed to get quite annoyed with this repeated discussion (see sidebar). It is interesting to note that, although a reservation for 121 bits NSAP addresses was made in RFC 1883, the use of ISO addresses was not a topic of discussion in the IPng working group meetings around December 1995.

"If this means that the Australian National Condom Board's requirement of having all Australian prophylactic manufacturers use 75 byte addresses conflicts with the desire of some of them to use our protocols, well, fine. Let them go and argue about that in an appropriate forum. This isn't the appropriate forum. This is the forum where people are discussing the engineering of IPv6, not whether the Australian National Antelope Commission will be happy with decisions made by the IESG. Lets get back to talking about the ENGINEERING of the IPv6, not the engineering of the ISO/OSI network stack, the politics of international standards bodies, or whether the New Caledonian Coastal Fishing Authority has recently passed a regulation requiring all tourists to be tatoood bright purple."

(Perry Metzger, in an email to the IPng mailing list, August 29, 1994)

During the months October to December of 1995, about the same topics were discussed as during the April and December working group meeting. Since the meeting attendees and mailing list contributors are essentially the same group of people, and the meeting agenda is discussed on the mailing list, this is only to be expected. A list of attendees to an interim meeting shows that the attendees are from government and military research laboratories (NIST, NRL), the computer industry (IBM, Sun, Digital) and from the networking industry (Ipsilon, Bellcore, Bay Networks).

In December 1998, the topics discussed on the mailing were again about the same as those on the IPng working group meetings. More interestingly, a starting discussion on quality of service, regarding the use of the Flow Label and the Traffic Class fields, was quickly diverted by pointing to the RSVP and Differentiated Services working groups. There was also much discussion about implementation issues. This led to a large list of things, outside of the core IPv6 protocol, that were still needed. Most of these were reflected in the 1998 working group meetings.

"Things are moving in the IETF but not as fast as the market needs IPv6 to happen [...] we will not appease the politically correct attitude, that has slowed up IPv6 thus far."

(Jim Bound, in an email to the IPng mailing list, December 4, 1998)

Articles from the tech press

A search in the technology press (CNET, IDG.NET, ZDNET) for articles on IPv6 shows a number of, mostly recent, articles on IPv6. Unfortunately, these articles say little on the standardisation process, but they do mention the current inhibition to wide-scale IPv6 deployment. The consensus here seems to be that demand for IPv6 will come chiefly for Asia-Pacific and developing countries, since only a fraction of the available IPv4 address space is allocated to these countries. Another important driver is the mobile communication sector, which is strong in Europe and also in the Asia-Pacific region. Since already 74% of the IPv4 address space is allocated in North America (Stanford University controls more IPv4 address space than the whole People's Republic of China), North America already has a large IPv4 installed base, and the mobile communications sector is divided there, demand for IPv6 products and services is not expected to come from this region first.

Critics and sceptics of the IPv6 transition say that using workarounds, like those mentioned in the Introduction, will be the preferred option of service providers over deployment of IPv6. On the other hand, most commentators seem to agree that "Deployment of IPv6 is not a matter of "if" but "when""[19].

4 Conclusions

The process that should lead to a successor to the currently ubiquitous Internet Protocol version 4 (IPv4) has been going for over a decade now. After an initial period which saw many diverging proposals, a direction was chosen in the beginning of 1995. The main standards were already done in 1998, but some parts still need some hard work. At the moment, the remaining problems are not with the new standard itself, but with the migration process from IPv4 to IPv6. There is standards-related work needed in that area also (like connecting "islands" of IPv6 networks in an ocean filled with IPv4), and working with IPv4 and IPv6 on one host computer at the same time. There are many details in the migration process that are still being discussed. A completely different kind of problem is formed by the chicken-end-egg problem, or, as Noel Chiappa, a long-time IPv6 critic, puts it: "It only makes people's life better if 90% or more switch to it, and because we can't [even] get from 1% to 2%, we're never going to get to 90%"[16].

In the previous section, three hypotheses have been formulated that might explain the lengthy IPv6 process. For the first hypothesis, that there has been a proliferation of requirements that caused the standardisation process to stall, there is not much evidence. The only piece of information hinting at this possibility is that a number of new topics have been introduced in the discussion at IPng meetings between 1995 and 1998. However, most of these topics relate to optional extensions, which are not part of the core IPv6 protocol, while others are related to problems already signalled in 1995. In 1998, remarks about e.g. the use of Quality of Service provisions of IPv6 is quickly relayed to other IETF working groups, so that the IPng

working group could concentrate on the core set of protocols. This would suggest a power of the working group to keep its own agenda in hand.

For the second hypothesis, that there has been little pressure on the IETF to "make haste", it is interesting to note that the first nervous comments about IPv6 maybe coming too late are made in the mailing list in 1998. The technology press adds to this that demand for IPv6 networking services is still mostly absent, especially in North America. Demand is expected to come from Europe and the Asia-Pacific region, where less IPv4 address space is available and where mobile communication is expected to surge far earlier than in the United States. Also, the European, the Japanese and the Korean governments have already pledged support for IPv6. For the period before 1998, no evidence was found for the feeling of time pressure by the IPng working group.

The third hypothesis, that "outsiders" (with respect to people traditionally associated with the IETF) have caused a break-up of the traditional IETF swiftness, finds support in two events. In both cases, the outsiders were related to the ISO set of networking protocols, and their influence was limited by the not-invented-here attitude of the IETF regulars.

To conclude, there is no convincing evidence for any of the three hypotheses. There is even some negative evidence for the first hypothesis, about proliferation of requirements of the new standard. More research on this subject, preferably including interviews with the individuals involved in the process, will be needed to say more on this.

5 References

- [1] J. Postel (editor), *Internet Protocol*, RFC 791, September 1981
- [2] A. Tan, *the LONG and windy ROAD*, September 2000, <http://itf.vlsm.org/road.html>
- [3] Internet Engineering Task Force, <http://www.ietf.org>
- [4] IPng Working Group, <http://www.ietf.org/html.charters/ipngwg-charter.html>
- [5] F. Kastenholz, C. Partridge, *Technical Criteria for Choosing IP: The Next Generation (IPng)*, RFC 1726, December 1994
- [6] S. Bradner, A. Mankin, *The Recommendation for the IP Next Generation Protocol*, RFC 1752, January 1995
- [7] S. Deering, R. Hinden, *Internet Protocol, Version 6 (IPv6) Specification*, RFC 1883, December 1995
- [8] S. Deering, R. Hinden, *Internet Protocol, Version 6 (IPv6) Specification*, RFC 2460, December 1999
- [9] R. Callon (ed.), *Minutes of the IPng Working Group (IPNGWG)*, April 1995, <http://playground.sun.com/ipv6/minutes/ipngwg-minutes-95apr.txt>
- [10] R. Hinden (ed.), *IPng W.G., Dallas IETF minutes*, December 1995, <http://playground.sun.com/ipv6/minutes/IPng-Meeting.Dec95.txt>
- [11] R. Hinden (ed.), *IPng Meeting, Chicago IETF minutes*, August 1998, <http://playground.sun.com/ipv6/minutes/ipng-meeting-aug98.txt>
- [12] R. Hinden (ed.), *IPng Working Group, December IETF minutes*, December 1998, <http://playground.sun.com/ipv6/minutes/ipng-wg-minutes-dec98.txt>
- [13] R. Hinden (ed.), *IPng Working Group Meeting minutes*, February 1999, <http://playground.sun.com/ipv6/minutes/ipng-minutes-feb99.txt>
- [14] R. Hinden (ed.), *IPng Meeting Minutes, San Diego IETF*, December 2000, <http://playground.sun.com/ipv6/minutes/ipng-minutes-dec2000.txt>
- [15] R. Hinden (ed.), *IPng Working Group, Minneapolis IETF meeting minutes*, March 2001, <http://playground.sun.com/ipv6/minutes/ipng-minutes-mar2001.txt>

- [16]D. Goodin, "Introducing a New Internet Protocol to Fix Traffic Problems Faces Criticism, Apathy", in *The Wall Street Journal*, May 14, 2001, <http://interactive.wsj.com/fr/emailthis/retrieve.cgi?id=SB98979192937941302.djm>
- [17]S. Taylor, "Cisco Launches Software for New Internet Devices", in *ZDNet Tech Infobase*, May 14, 2001, <http://cma.zdnet.com/>
- [18]E. Krapf, "Whatever Happened To IPv6?", in *Business Communications Review*, January 1, 2001, <http://cma.zdnet.com/>
- [19]N. Petrosky, "IPv6: This is your wake-up call", in *NetworkWorldFusion*, May 14, 2001, <http://www.nwfusion.com/columnists/2001/0514petrosky.html>

Appendix: IPng Working Group meeting discussion topics

Topic	1995	1998	2001
Addressing Architecture	X		X
Mobility for IPv6 (Mobile-IP WG)	X	X	
Explicit Routing Header	X		
Provider-based Unicast Addressing	X		X
ICMP for IPv6	X		
DNS extensions for IPv6	X	X	
Neighbour Discovery	X	X	
Security	X	X	
Unicast Address Allocation Architecture	X		X
Multihoming	X		X
Path MTU Discovery	X	X	
Anycast addressing for hosts	X		
Socket API	X	X	
Tunneling	X	X	X
Higher-level protocols like FTP	X		
Routing protocols: RIP, OSPF, IDRP, BGPng, multicast routing	X		
Management Information Bases (MIBs)	X		X
Renumbering		X	
Link-local/Site-local addressing in relation to address selection, mobility etc.		X	X
Reserved anycast addresses		X	
IPv6 over IPv4 without explicit tunnels		X	
Jumbograms		X	
Multicast Listener Discovery Protocol		X	X
Advanced Socket API		X	
Addressing: Sub-TLA assignments		X	
Router renumbering		X	
Site Prefixes in Neighbour Discovery		X	X
IPv6 TCP & Anycast		X	
RPC, NFS over IPv6		X	
Tunnel Broker		X	
TLA/NLA Assignments Rules		X	
IPv6 Header Compression		X	
DHCPv6		X	
Use of Flow Labels		X	X
Secure Neighbour Discovery&Autoconfiguration		X	
IPv6 for PDAs and mobile phones		X	
Plug&Play: Automatic finding of DNS servers, services like printers, default router		X	
Billing/Accounting opportunities		X	
E.164-in-IPv6 addressing		X	

Table 2: Points for discussion in IPng Working Group meetings in 1995, 1998 and 2001