CHAPTER 1

Overview

During the years between the end of the second millennium and the beginning of the third one, computer networks will benefit from the availability of many new technologies, including ATM, Gigabit Ethernet, and virtual LANs. The organization of the Internet and of Intranets will have a strong evolution thanks to the adoption of the new IPv6 protocol.

But what is IPv6? IPv6 is the new version of the *IP protocol* (Internet Protocol) on which the Internet and many Intranets are based. The work for IPv6 standardization began in 1991, and the main part was completed within 1996 with the publication of *RFCs* (Requests For Comments), standards that exactly define IPv6. During the standardization phase, this new protocol was indicated also by the terms *IPng* (IP new generation) and *IPv7*. What happened to IPv5? It lost the race, and therefore everyone agreed not to use that version number.

This book moves from the author's firm belief that, in the interim, IP will be the only layer 3 protocol to survive.

This didactic text provides a global overview of the protocol organization, of its functions, and of problems related to its adoption "in the field." In this sense, this book cannot and will not replace standard RFCs, to which readers must refer to resolve their doubts if they want to get into further details or they must deal with the design of IPv6-based plants, products, networks, and so on.

1.1 Why IPv6?

The answer is simple: "The Internet is becoming a victim of its own success." Probably many of you have heard this sentence repeated many times lately, but what does it really mean?

Ordinary users see the Internet through its applications they use daily for their work—from electronic mail, which has become user-friendly thanks to application software such as Eudora and Pegasus, to the navigation on WWW servers with powerful browsers such as Netscape or Microsoft Explorer, which today are frequently enriched with Java applets. In general, users have had a great deal of success with all Internet applications, even the more simple ones such as FTP or Telnet, and many companies have decided to reorganize their networks on the Internet model by creating Intranets.

The worldwide success of the Internet and of Intranets keeps pace with the success of the network architecture called *Internet Protocol Suite*, best known as TCP/IP, on which they are based.

In particular, the present IP protocol (Internet Protocol) is a protocol standardized in 1981 by RFC 791[1]; therefore, this protocol is a little dated even if it is a cornerstone of the architecture. To avoid confusion, in the following text we will indicate the present IP protocol that has version number 4 with the acronym *IPv4*, the new protocol with the acronym *IPv6,* and we will simply use *IP* to indicate what is common to both versions.

IP handles the decoupling of applications from transmission networks; that is, it enables users to use their preferred applications independently from the underlying network technology (see Figure 1-1).

Moreover, IP allows users to use different technologies in different parts of the network—for example, LANs (Ethernet, Token Ring, FDDI) inside buildings and frame relay or ATM public services for the geographic part of the same network.

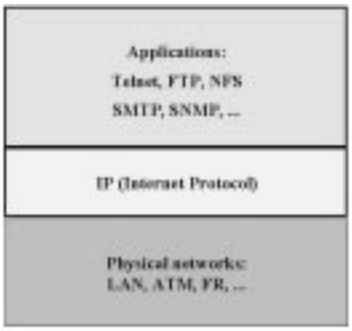

Figure 1-1
Internet Protocol (IP)

IPv4 achieves this result by providing a service with the following main characteristics:

- *Universal addressing:* Each IPv4 network interface has a unique worldwide address with 32 bits.
- *Best effort:* IPv4 performs its best effort to deliver packets, but it doesn't guarantee anything at the upper layer, neither in terms of percentage of delivered packets nor in terms of time used to execute the delivery. In short, IPv4 doesn't have a built-in concept of *Quality of Service* (QoS).

These two characteristics, which have been points of strength for IPv4 up to now, risk becoming its main limits and forcing the introduction of IPv6. Let's look at the reasons.

1.1.1 Why a New Address Scheme?

We have already seen that IPv4 addresses take up 32 bits, which means that in total about 4 billion addresses are available and, because 4 billion computers don't exist in the world, understanding the reasons that the Internet is running out of addresses is not immediately apparent. We must search for the reasons in the IPv4 address structure and in assignment procedures, which cause a significant number of assigned addresses to be unused.

In fact, IPv4 addresses are not assigned one by one (a procedure clearly impossible for organizational reasons), but by "networks." Networks belong to three different classes:

- Class A: 128 available networks, each one with about 16 million addresses

- Class B: About 16,000 available networks, each one with about 65,000 addresses
- Class C: About 2 million available networks, each one with 254 addresses

In January 1996, 92 class A networks, 5655 class B networks, and 87,924 class C networks were assigned. This data shows that the main problem is related to class B networks, which, for their intermediate size, are more suitable to be assigned to organizations. In fact, class A networks are too wide, and only 36 are left to be assigned, whereas class C networks are too small. Table 1-1 shows the growth trend of networks and addresses.

The problem of IPv4 address exhaustion was realized in 1991. In that year, the requests for address assignments began to grow more rapidly than any expectations. It was a historic moment when the Internet became the only network for everybody. And when we say *everybody,* we really mean everybody: public and private companies, government and private administrations, universities and research centers, and above all, private citizens. This use was made possible by ISPs (Internet Service Providers)

Table 1-1

Growth in time of networks and IPv4 addresses

Date	Host	Networks of Class:		
		A	B	C
Jan 97	16,146,000			
Jun 96	12,881,000			
Jan 96	9,472,000	92	5655	87,924
Jul 95	6,642,000	91	5390	56,057
Jan 95	4,852,000	91	4979	34,340
Oct 94	3,864,000	93	4831	32,098
Jul 94	3,212,000	89	4493	20,268
Jan 94	2,217,000	74	4043	16,422
Oct 93	2,056,000	69	3849	12,615
Jul 93	1,776,000	67	3728	9,972
Apr 93	1,486,000	58	3409	6,255
Jan 93	1,313,000	54	3206	4,998

that provide low-cost connections to the Internet through telephone lines first by using modems and, more recently, ISDN access. A further turning point is very recent: the introduction of xDSL and "cable modems" to provide all domestic users with high-speed connections to the Internet (faster than 1 Mbps).

In 1991, forecasts were that class B addresses would be used up within 1994. To face this dramatic forecast and to leave a reasonable amount of time for the development and the migration to IPv6, the IETF (Internet Engineering Task Force), the committee responsible for technical decisions for IP and for the Internet, decided to assign not only class B networks, but also blocks of class C "adjacent" networks. For example, an organization with 100 computers with a growth forecast to 500 computers could be assigned, instead of a class B network, a block of four class C networks for a total of about 1000 addresses.

This new and more conservative policy of address assignment moves forward the moment in which IPv4 addresses will be exhausted: Some very uncertain forecasts identify a date between 2005 and 2015.

There is no rose without a thorn, as an old saying goes, and also this addressing scheme immediately generates problems on routers that are forced to maintain routing information for each network. In fact, if an organization is assigned a class B network, routers must have only one routing entry, but if it is assigned 16 class C networks, routers must have 16 different routing entries, using 16 times more memory for routing tables. To avoid this problem, the CIDR (Classless InterDomain Routing)[2] was introduced in 1992, which in substance means that the concept of network class at the routing table level is eliminated.

In the end, the suggestion is that all Intranets use the same addresses, and to this purpose the RFC 1597[3] was issued, later replaced by the RFC 1918[4], assigning Intranets a class A network (the 10.0.0.0) and some class B and C networks.

At this point, it should be clear that IPv6 needs a new addressing scheme with the following characteristics:

- A higher number of bits so that the addressing space is not subject to further exhaustion
- A more flexible hierarchical organization of addresses that doesn't use the concept of classes, but the CIDR mechanism
- A scheme for address assignment aimed to minimize the size of routing tables on routers and to increase the CIDR performance
- Global addresses for the Internet and local addresses for Intranets

1.1.2 Best Effort: Is It Enough?

IPv4 is a connectionless protocol. This means that it transmits each packet independently from other ones, specifying in the packet header IPv4 addresses of the source and of the destination. The packet is neither marked as belonging to a flow or to a connection, nor numbered in any way. Therefore, it is neither possible to correct errors at this level nor to understand whether a packet has been delivered, or if so, what was the delivery time. This kind of service is called "best effort" because every IPv4 node performs at its best to deliver the packet in the minimum time, but it cannot guarantee if and when the delivery will happen.

Best effort connectionless protocols can be implemented easily and have a limited and constant overhead. These characteristics allowed IPv4 to become popular—and eventually the only surviving layer 3 protocol.

Nevertheless, the availability of new high-speed ATM networks guaranteeing the QoS[5], on the one hand, and the need to develop new multimedia applications requiring a guaranteed QoS, on the other hand, have led to discussions of whether "best effort" choice is still to be considered the best one for IPv6.

The IETF has already recognized the lack of the concept of QoS as a limit of IP, and it has developed an additional protocol, called RSVP (Resource reSerVation Protocol)[6], to allocate resources on routers and make them suitable to guarantee the QoS for IPv4-based applications that explicitly require a given QoS through RSVP.

IPv6, while remaining faithful to the IPv4 connectionless origin, introduces the concept of flow as a better integration mechanism toward QoS concepts and with RSVP.

1.2 Requirements to Be Met by IPv6

Up to now, we have discussed reasons to switch from IPv4 to IPv6, and we have caught a glimpse of some characteristics that differentiate IPv6 from IPv4. The question to be answered now is: Which characteristics do we want to maintain, which ones do we want to eliminate, and which new ones do we want to introduce?

A risk that the IETF has always taken into consideration is the "second generation syndrome," which consists of adding everything that users ask with the risk of obtaining a slow, not manageable, and useless protocol.

Let's inspect the main expectations that emerged about IPv6[7].

1.2.1 An Address Space to Last Forever

The expectation here mainly depends on what we mean by the term *forever*. A proposal could be to have an IPv6 address for every potential Internet user. We can estimate that the world population will reach 10 billion people and assume that each person will have more than one computer because, in the future, home appliances, electro-medical devices, and electrical devices in general will be computers. Today, we already have available domestic lighting systems in which lamps have an address and are turned on and off by messages sent by switches on a service bus. In the future, Internet users might want to order from outside their homes that an oven begin to cook a turkey, or to receive a message from their home alarms to detect a possible intrusion, or to control their Internet browsers using remote-controlled video cameras. The examples are diverse; cellular telephones with Java terminals inside already appear on the market. An estimate of 256 IPv6 addresses for each planet inhabitant is not unrealistic.

A more drastic proposal is to try to estimate the number of IPv6 addresses based on the number of atoms in the universe, keeping in mind that you only need about an atom to build a computer. But, be careful not to exaggerate; in fact, having more addresses means a greater length of IPv6 address fields, and because both the source and the destination address must be transported within each IPv6 packet header, this means more overhead.

On the other hand, everybody agrees to define an addressing space that is not subject to exhaustion in the future.

Besides the number of addresses to be assigned, considering the efficiency of the assignment scheme is also important. An accurate study by Christian Huitema[8] proposes to define the efficiency of address assignment H as the ratio between the logarithm in base 10 of the number of used addresses and the address bits number.

$$H = \frac{\log_{10}(\text{address number})}{\text{bits number}}$$

In a scheme with a maximum efficiency rate, all addresses are used; therefore, H is equal to the base 10 logarithm of 2 (that is, H = 0.301). An analysis of real addressing schemes shows that H varies between 0.22 and 0.26.

The final decision is to predict one million billion networked computers (10^{15}) that, with H equal to 0.22 (the worst case), require 68-bit addresses. Because the address, for implementation reasons, must be a multiple of 32 bits, it has been opted for having the IPv6 address on 128 bits (that is, 16 bytes or 4 words of 32 bits).

1.2.2 Multicast and Anycast Addresses

Besides Layer 3 unicast addresses (described previously), IPv4 also utilizes multicast or class D addresses for applications that require group communications such as video conferencing on the Internet. The concept of multicast addresses is also handled in IPv6.

IPv6 also introduces a new type of address called *anycast*. These addresses also are group addresses in which the only member of the group to respond is the "closest" to the source. The use of anycast addresses is potentially very interesting because the closest router, the closest name server, or time server can be accessed by an anycast address.

1.2.3 To Unify Intranets and the Internet

IPv6 must provide a unified addressing scheme for the Internet and for Intranets, overcoming temporary IPv4 solutions (RFC 1597[3] and RFC 1918[4]). For this purpose, besides global addresses, site addresses and link local addresses also have been developed. Site addresses should be used for network nodes inside Intranets, whereas link local addresses are used to identify nodes attached to a single link (small networks without a router).

Lastly, addresses with embedded IPv4, OSI NSAP, and Novell IPX addresses have been developed.

1.2.4 Using LANs Better

When IPv4 operates on a LAN, it frequently needs to determine the relationship between an IPv4 address and a MAC address, and vice versa. IPv4 performs this function through an auxiliary protocol called *ARP* (Address Resolution Protocol)[9] that utilizes broadcast MAC layer transmissions. A broadcast packet is received by all stations and causes an interruption on all stations, including those not using the IP protocol. This ineffectiveness must be corrected in IPv6 by using a "neighbor discovery" method on LAN more efficient than ARP and utilizing multicast, not broadcast, transmissions. In fact, a station can determine at the network adapter level which multicast to receive, while it is obliged to receive all broadcasts.

1.2.5 Security

The security in IPv4 is today managed through particular routers or computers performing the role of *firewalls*. They cannot solve intrinsic IPv4 security problems, but they can counterbalance many computers' operating system weaknesses and the superficial management of security that frequently exists at a single computer level.

IPv6 is not necessarily requested to improve the security state of the art, but it will not make the situation worse. As a matter of fact, the IETF defined a series of encryption and authentication procedures that will be available in the IPv6 protocol in the beginning. These procedures will also be implemented in a compatible way in IPv4.

Moreover, IPv6 has a careful management of *Source Routing*, that is, of the possibility to determine at source station level the path to be followed by an IP packet. This function, already available in IPv4 but not always implemented or active, is frequently exploited by hackers to try to bypass firewalls.

Many network administrators will undoubtedly find in the availability of standard security procedures one of the main reasons for migrating to IPv6.

1.2.6 Routing

Routing is clearly one of the central themes in the design of a protocol expected to route packets on the future Internet. If we consider IPv4 routing as a starting point, we can see that routing tables of Internet routers tend to explode. In fact, if the CIDR is not used, every single network must be announced by an entry in routing tables. The CIDR introduction[2] allows us to announce a block of networks with contiguous addresses (for example, 195.1.4.0, 195.1.5.0, 195.1.6.0, and 195.1.7.0) as a unique entry by specifying how many bits must be considered as significant (in our example, 195.1.4.0/22, which is each network with the first 22 bits equal to 195.1.4.0).

In any case, the CIDR can do little if it is not connected to the address assignment. In fact, if addresses are assigned to ISPs (Internet Service Providers) and by them to users, the CIDR works properly because, from a theoretical point of view, all addresses of a single ISP can be announced by a unique entry. We can think of a form of hierarchical routing accompanied also by a hierarchical kind of address assignment bound to the network topology. At the root of the hierarchical tree, we can think of an address assignment by continents; then within a continent, an assignment by ISPs; then by organizations; and eventually by networks within organizations. This model minimizes tables on routers, allowing the CIDR to aggregate addresses first by user, then by ISP, and eventually by continent, but this model has a big limit: The users don't have any more addresses permanently assigned to them.

If we consider how the IPv4 address assignment is managed nowadays, an organization can contact authorities such as INTERNIC (Northern America), APNIC (Asia and Pacific) and RIPE-NCC (Europe) to obtain addresses that the organization will use independently from the ISP it will be connected to. This way, the organization can change ISPs without changing addresses. With IPv6, when an organization changes ISPs, it necessarily must change addresses. An organization may even have to change addresses because two ISPs have merged or separated; therefore, the organization must change addresses even if it doesn't want to.

The address assignment model based on the network topology is acceptable in IPv6 only if autoconfiguration mechanisms (plug and play) are available (that is, networks dynamically assign addresses to stations).

So far, we have talked about computation of routing tables used for default routing toward a given destination. IPv6 also addresses the possibility of having policy routing and QoS (in this context called *ToS*, or *Type of Service*). An example of routing based on a particular policy is one that

determines the transmission of packets to a given destination on a path determined also by the source address (this was impossible in Ipv4).

The IPv6 routing must also provide good support for mobility—for example, to those users who, by means of a portable PC and a cellular phone, can connect themselves to the Internet in different places.

1.2.7 A Good Support for ATM

The great industrial effort related to the development of *ATM* (Asynchronous Transfer Mode)[5] will make this technology one of the most important actors in future wide area and local area networks. IPv6 designers, well aware of this fact, tried to improve the support of ATM in IPv6. But what are ATM's peculiarities? ATM is an NBMA (Non-Broadcast Multiple Access) network, and it guarantees the QoS.

An NBMA network[10] is a multipoint access network that doesn't provide a simple mechanism to transmit a packet to all other stations. IPv4 has been designed to work either on point-to-point channels that have only two endpoints or on local networks that have multiple access, but where a packet transmission to a single station or to all stations has exactly the same cost. Other NBMA networks are, for example, X.25 and Frame Relay (if equipped with signaling), but the need to provide a good IP support on NBMA networks emerged only with ATM because of the role that this technology will play in the future.

Guaranteeing the QoS means associating to each data flow a given set of quality requirements. For example, if the data flow has been generated by a file transfer, that the loss rate is equal to zero is very important, whereas the delay to which packets are subject along the path is irrelevant. If the data flow is generated by an audio or video source, a certain rate of loss of data can be tolerated (we can understand audio and video signals also if uncompleted), but guaranteeing limited and less variable delays from a packet to another is fundamental.

We must also remember that the QoS can be used only if it is requested by applications, an action that today's applications don't perform. We need to foresee that applications request the QoS through a protocol like RSVP[6] (see Section 1.2.2) and that this one, by jointly operating with IPv6, transforms the QoS request into a QoS request for the ATM network (see Figure 1-2).

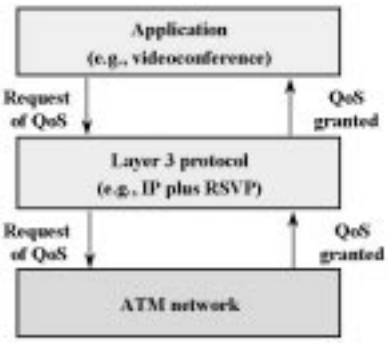

Figure 1-2
Handling of QoS requests

1.2.8 The Concept of Flow

To simplify the implementation of IPv6 on ATM and the QoS management, we need to introduce the concept of *flow*. A flow is a sequence of packets in some way correlated (for example, because they have been generated by the same application) and that therefore must be treated coherently by the IP layer. Packets belong to the same flow on the basis of parameters like the source address, the destination address, the QoS, the accounting, the authentication, and the security.

No relationships exist between the concept of flow and other concepts such as TCP connection; for example, a flow can contain several TCP connections. Moreover, we must emphasize that the introduction of the concept of flow occurs on a protocol that is and remains connectionless (also frequently called a *datagram*); therefore, flows do not have the same purposes of connection-oriented protocols—for example, correction of errors. In general, a flow can have as its destination either a single station or a group of stations; therefore, we can have either unicast or multicast flows.

After the concept of flow has been introduced, we can introduce the flow label concept by which we will mark packets or datagrams by reserving a special field in the IPv6 header. In this way, IPv6 has the possibility, at the moment it receives a packet, to know to which flow it belongs by examining its flow label and, as a result, to know the packet needs in terms of QoS.

1.2.9 Priorities

Even if an application doesn't request a QoS, differentiating the traffic generated by principal applications as a function of their real-time requirements is possible. For this purpose, a 4-bit "priority" field has been

introduced in the IPv6 header to differentiate 16 potential traffic priorities. Up to now, priorities have been defined for news, e-mail, FTP, NFS, Telnet, X, routing, and SNMP protocols.

1.2.10 Plug and Play

In Section 1.3.1, we saw how IPv6 needs autoconfiguration (or plug and play) mechanisms to manage addresses that can change in the long run. Moreover, manual management is inconvenient because an IPv6 address requires that 32 hexadecimal digits be written (for example, **`FEDC:BA98:1234:5678:0BCA:9987:0102:1230`**).

The *DHCP* (Dynamic Host Configuration Protocol)[11], available on some IPv4 implementations, has been considered a good starting point. The idea is to develop a DHCPv6 protocol that allows the automatic configuration of hosts and subnetworks, the learning of default routers, and through an interaction with the DNS (Domain Name Service)[12], also an automatic configuration of host names.

The implementation of the DHCPv6 on all IPv6 hosts will allow network administrators to reconfigure addresses by operating on the primary DHCPv6 server.

1.2.11 Mobility

As we already mentioned, an increasing number of Internet users don't work at their office desks anymore but work while traveling. Mobile users are usually equipped with portable PCs with the PCMCIA network card, which connects them to a cellular telephone or to a public network via radio.

IPv4 doesn't provide any support for mobility. In fact, every computer has a fixed address that belongs to a network. If the computer is connected to a different network, packets sent to it continue to reach the original network, and there they are lost.

Clearly, providing support for mobility is a main requirement for IPv6: It has been estimated that, in Northern America, there will be from 20 to 40 million mobile users in 2007. Also, this requirement is one of the more complex to be met, as it has to deal with a range of problems, starting from those related to radio transmission (reliability, roaming, hand-off) to those related to IP protocols (identification, addressing, configuration, routing) to security problems.

The solution that is taking shape predicts that mobile users will have two addresses: the first one "permanent" on their organization's network and the second one "dynamic" depending on the point from which they are connected in a given moment. The organization's firewall, when the users are traveling, acts as "proxy" for the permanent address and creates a safe tunnel toward the dynamic address.

1.2.12 Transition from IPv4 to IPv6

Many users will consider the transition to IPv6 as something they must resign themselves to so that they can obtain the potential advantages discussed previously. But people, like me, who have experienced other transitions know that, even if such transitions are well planned, they can easily end up as a "blood bath." Changing the network software is similar to changing the operation system version: This step potentially brings forward some incompatibilities and causes the need to update both the hardware and the software.

The IETF decided to design a migration strategy based on a "dual-stack" approach, but this approach will be a field in which computer and network vendors will fight strongly to simplify users' lives and to win market share. In fact, very few users will be able to migrate at a given moment; many organizations will have a transition period lasting months or even years, during which IPv6 must coexist with IPv4.

For this reason, the IETF decided that IPv4 and IPv6 will be two different protocols with two corresponding and separated protocol stacks. When a station receives a frame from its local network, the *Protocol Type* allows it to distinguish whether the frame contains an IPv4 or an IPv6 packet, with the same mechanisms that allow it to distinguish between IPv4 and Decnet packets today. In fact, we know that IPv4 packets have a protocol type equal to 0800H (800 Hexadecimal), and IPv6 packets have a protocol type equal to 86DDH.

Therefore, the first field of IPv4 and of IPv6 packets, representing the protocol version (that can assume values 4 or 6), will remain unused because the IPv4 stack will receive only IPv4 packets and the IPv6 stack will receive only IPv6 packets.

One of the critical steps in the transition will be the parallel management of IPv4 and IPv6 addresses. A timely updating of DNS servers will be necessary, followed by the updating of DHCP servers. A dual-stack station will use the IPv4 address (32 bits wide) to communicate with other IPv4 stations, and it will use the IPv6 address (128 bits wide) to communicate with other IPv6 stations.

For this approach to be successful, IPv6 islands must be interconnected. This connection will be implemented through a series of tunnels on the Internet, and therefore on IPv4, that will form a layered network called *6-Bone*. This approach is based on the positive experience of *Mbone*, the network used for video conferencing on the Internet, that has been successfully implemented following the same philosophy.

6-Bone will grow and some islands will directly interconnect using IPv6, without needing tunnels. An increasing number of machines will communicate by using IPv6; then the end of IPv4 will arrive, when all computers running only the IPv4 protocol stack will lose their direct global connectivity to the Internet.

1.3 Choice Criteria

The need to meet all these requirements reveals how difficult the choice of the new IPv6 has been, because this protocol will be entrusted with the destiny of the Internet and Intranets. The previously listed requirements are joined by another one to maintain the critical router loop simply. The *critical router loop* is the set of code lines that route most packets, all those packets that don't have particular requests apart from reaching the destination. The critical router loop determines the router's performance more than any other part of the code, and a careless addition of all the new requested and previously mentioned functions will complicate the situation too much.

For this reason, IPv6 designers Steven Deering and Robert Hinden decided to take to themselves a famous maxim by Antoine de Saint-Exupery, the author of *The Little Prince*, a nice book that I suggest everybody read, about architectural simplicity:

> The architectural simplicity
>
> In each thing, you reach the perfection, not when there is nothing left to add, but when there is nothing left to take off.
>
> *Antoine de Saint-Exupery*

The result is a protocol with an extremely pure design and a small header with few fields. In fact, the IPv4 header (see Figure 1-3) consists of 24 bytes, 8 of which are used for IPv4 addresses and the remaining 16 bytes by 12 additional fields.

The IPv6 header (see Figure 1-4) has only 40 bytes, 32 of which are used for IPv6 addresses and the remaining 8 bytes by 6 additional fields.

And what about all the fields needed to implement many new additional functions? They have been inserted in various *extension headers*

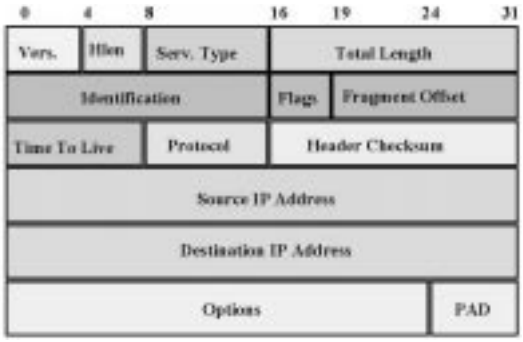

Figure 1-3
The IPv4 header

Figure 1-4
The IPv6 header

that are present only if the function is effectively requested. In this way, most packets pass very quickly through critical router loops, and only packets with particular requests receive a more sophisticated treatment that provides for the extension header's analysis. In any case, many extension headers have "end-to-end" functions; therefore, they don't need to be processed by routers, but only by source and destination nodes. (A typical example is represented by the encryption extension header.)

1.4 The Path Toward Standardization

The path toward standardization formally began in 1992, when the IETF, during a meeting in Boston, issued a "call for proposal" for IPv6 and many working groups were created.

The main proposals for IPv6 are described in the following subsections.

1.4.1 TUBA

The proposal known as TUBA (TCP and UDP over Bigger Addresses)[13] suggested the adoption of the ISO/OSI 8473 CLNP protocol to replace IPv4, trying in this way to create a fusion *in extremis* between the OSI world and the Internet world. This solution would have allowed users to have at their disposal OSI NSAP 20-byte addresses and a common platform on which OSI transport protocols, such as TP4 and the cited TCP and UDP, could be used.

The main censure made against CLNP by the Internet world was that it had been copied 10 years before from IPv4 by introducing some depreciatory modifications.

Supporters of the TUBA proposal, in the first two years of discussions, remained faithful to the original CLNP project, refusing to introduce innovative aspects such as multicasting, mobility, and QoS for reasons of incompatibility with the OSI installed base (of secondary importance). This stubbornness brought about the failure of the TUBA proposal, later followed by a general failure of the OSI CLNP.

1.4.2 IPv7, TP/IX, CATNIP

In 1992, Robert Ullmann advanced the proposal of a new IP protocol called IPv7. The proposal was re-elaborated in 1993 and assumed the name of TP/IX to indicate the will to change both the IP protocol and the TCP protocol at the same time. The proposal contained interesting ideas about speed packet processing and a new routing protocol called RAP. In 1994, the proposal had a further evolution, trying to define a unique format for IP, CLNP, and IPX packets, and assumed the new name of CATNIP[14]. CATNIP would have been a common platform supporting several transport protocols such as OSI/TP4, TCP, UDP, and SPX. Layer 3 addresses adopted by CATNIP were of OSI/NSAP type.

1.4.3 IP in IP, IPAE

IP in IP was a proposal made in 1992, designed to use two IPv4 layers to limit the address shortage at the Internet level: a layer to implement a worldwide backbone and a second layer within limited areas. In 1993, the proposal was developed further and was called IPAE (IP Address Encapsulation) and accepted as a transition solution toward SIP.

1.4.4 SIP

SIP (Simple IP) was proposed by Steve Deering in November 1992. It was based on the idea of bringing IP addresses to 64 bits and to eliminate some obsolete IPv4 details. This proposal was immediately accepted by many companies who appreciated its simplicity.

1.4.5 PIP

PIP (Paul's Internet Protocol), a proposal by Paul Francis, introduced significant innovations on the front of routing by allowing an efficient policy routing and mobility implementation. In September 1993, PIP merged with SIP, thus creating SIPP.

1.4.6 SIPP

SIPP (Simple IP Plus)[15] tried to combine the implementation simplicity of SIP and the routing flexibility of PIP. SIPP was designed to work efficiently on high-performance networks, such as ATM, but also on low-performance networks, such as wireless networks. SIPP has a small size header and 64-bit addresses.

The header coding is particularly emphasized. With SIPP, the header can be efficiently elaborated by routers and can be extended to insert new options in the future.

1.5 The Evaluation

A comparative evaluation of the last three proposals (CATNIP, SIPP, and TUBA) brought about the results shown in Table 1-2.

Table 1-2

Comparative analysis of three proposals for IPv6

	CATNIP	SIPP	TUBA
Complete specification	no	yes	mostly
Simplicity	no	no	no
Scale	yes	yes	yes
Topological flexibility	yes	yes	yes
Performance	mixed	mixed	mixed
Robust service	mixed	mixed	yes
Transition mechanisms	mixed	no	mixed
Media independence	yes	yes	yes
Connectionless service (datagram)	yes	yes	yes
Configuration simplicity	unknown	mixed	mixed
Security	unknown	yes	mixed
Name uniqueness	mixed	mixed	mixed
Standards access	yes	yes	mixed
Multicast support	unknown	yes	mixed
Extensibility	unknown	mixed	mixed
Availability of service classes	unknown	yes	mixed
Mobility support	unknown	mixed	mixed
Control protocol	unknown	yes	mixed
Tunneling support	unknown	yes	mixed

1.6 The Final Decision

The decision made in June 1994 was to adopt SIPP as a base for IPv6 with the modification of the address length from 64 to 128 bits.

1.7 Conclusion

The point of no return has been passed, a new IP protocol is at last a standard, and it will be a main actor in our future. Some competitors have been defeated, and among them the worst defeat was to OSI CLNP. But

now it is time to forget *ifs* and *buts* and to begin to work on these new standards. Currently, RFCs from **17** to **36** are already available.

REFERENCES

[1]J. Postel, *RFC 791: Internet Protocol*, September 1981.

[2]V. Fuller, T. Li, J. Yu, K. Varadhan, *RFC 1519: Classless Inter-Domain Routing (CIDR): An Address Assignment and Aggregation Strategy*, September 1993.

[3]Y. Rekhter, B. Moskowitz, D. Karrenberg, G. de Groot, *RFC 1597: Address Allocation for Private Internets*, March 1994.

[4]Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, E. Lear, *RFC 1918: Address Allocation for Private Internets*, February 1996.

[5]Uyless Black, *ATM: Foundation for Broadband Networks*, Prentice Hall, 1995.

[6]B. Braden, L. Zhang, D. Estrin, S. Herzog, S. Jamin, *RSVP: Resource ReSerVation Protocol (RSVP)— Version 1 Functional Specification*, Work in progress, January 1996.

[7]S.O. Bradner, A. Mankin, *IPng: Internet Protocol Next Generation*, Addison-Wesley, 1995.

[8]C. Huitema, *IPv6: The New Internet Protocol*, Prentice-Hall, 1996.

[9]D.C. Plummer, *RFC 826: Ethernet Address Resolution Protocol: On converting network protocol addresses to 48 bit Ethernet address for transmission on Ethernet hardware*, November 1982.

[10]J. Heinanen, R. Govindan, *RFC 1735: NBMA Address Resolution Protocol (NARP)*, December 1994.

[11]R. Droms, *RFC 1541: Dynamic Host Configuration Protocol*, October 1993.

[12]P.V. Mockapetris, *RFC 1035: Domain names—implementation and specification*, November 1987.

[13]R. Callon, *RFC 1347: TCP and UDP with Bigger Addresses (TUBA), A Simple Proposal for Internet Addressing and Routing*, June 1992.

[14]M. McGovern, R. Ullmann, *RFC 1707: CATNIP: Common Architecture for the Internet*, October 1994.

[15]R. Hinden, *RFC 1710: Simple Internet Protocol Plus White Paper*, October 1994.

[16]S. Bradner, A. Mankin, *RFC 1752: The Recommendation for the IP Next Generation Protocol*, January 1995.

[17]C. Partridge, *RFC 1809: Using the Flow Label Field in IPv6*, June 1995.

[18]IAB, IESG, *RFC 1881: IPv6 Address Allocation Management*, December 1995.

[19]S. Deering, R. Hinden, *RFC 1883: Internet Protocol, Version 6 (IPv6) Specification*, December 1995.

[20]R. Hinden, S. Deering, *RFC 1884: IP Version 6 Addressing Architecture*, December 1995.

[21]A. Conta, S. Deering, *RFC 1885: Internet Control Message Protocol (ICMPv6)*, December 1995.

[22]S. Thomson, C. Huitema, *RFC 1886: DNS Extensions to support IP version 6*, December 1995.

[23]Y. Rekhter, T. Li, *RFC 1887; An Architecture for IPv6 Unicast Address Allocation*, December 1995.

[24]R. Hinden, J. Postel, *RFC 1897: IPv6 Testing Address Allocation*, January 1996.

[25]R. Elz, *RFC 1924: A Compact Representation of IPv6 Addresses*, April 1996.

[26]R. Gilligan, E. Nordmar, *RFC 1933: Transition Mechanisms for IPv6 Hosts and Routers*, April 1996.

[27]T. Narten, E. Nordmark, W. Simpson, *RFC 1970: Neighbor Discovery for IP Version 6 (IPv6)*, August 1996.

[28]S. Thomson, T. Narten, *RFC 1971: IPv6 Stateless Address Autoconfiguration*, August 1996.

[29]M. Crawford, *RFC 1972: A Method for the Transmission of IPv6 Packets over Ethernet Networks*, August 1996.

[30]M. Crawford, *RFC 2019: Transmission of IPv6 Packets Over FDDI*, October 1996.

[31]D. Haskin, E. Allen, *RFC 2023: IP Version 6 over PPP*, October 1996.

[32]D. Mills, *RFC 2030: Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI*, October 1996.

[33]Y. Rekhter, P. Lothberg, R. Hinden, S. Deering, J. Postel, *RFC 2073: An IPv6 Provider-Based Unicast Address Format,* January 1997.

[34]G. Malkin, R. Minnear, *RFC 2080: RIPng for IPv6*, January 1997.

[35]R. Gilligan, S. Thomson, J. Bound, W. Stevens, *RFC 2133: Basic Socket Interface Extensions for IPv6,* April 1997.

[36]D. Borman, *RFC 2147: TCP and UDP over IPv6 Jumbograms,* May 1997.