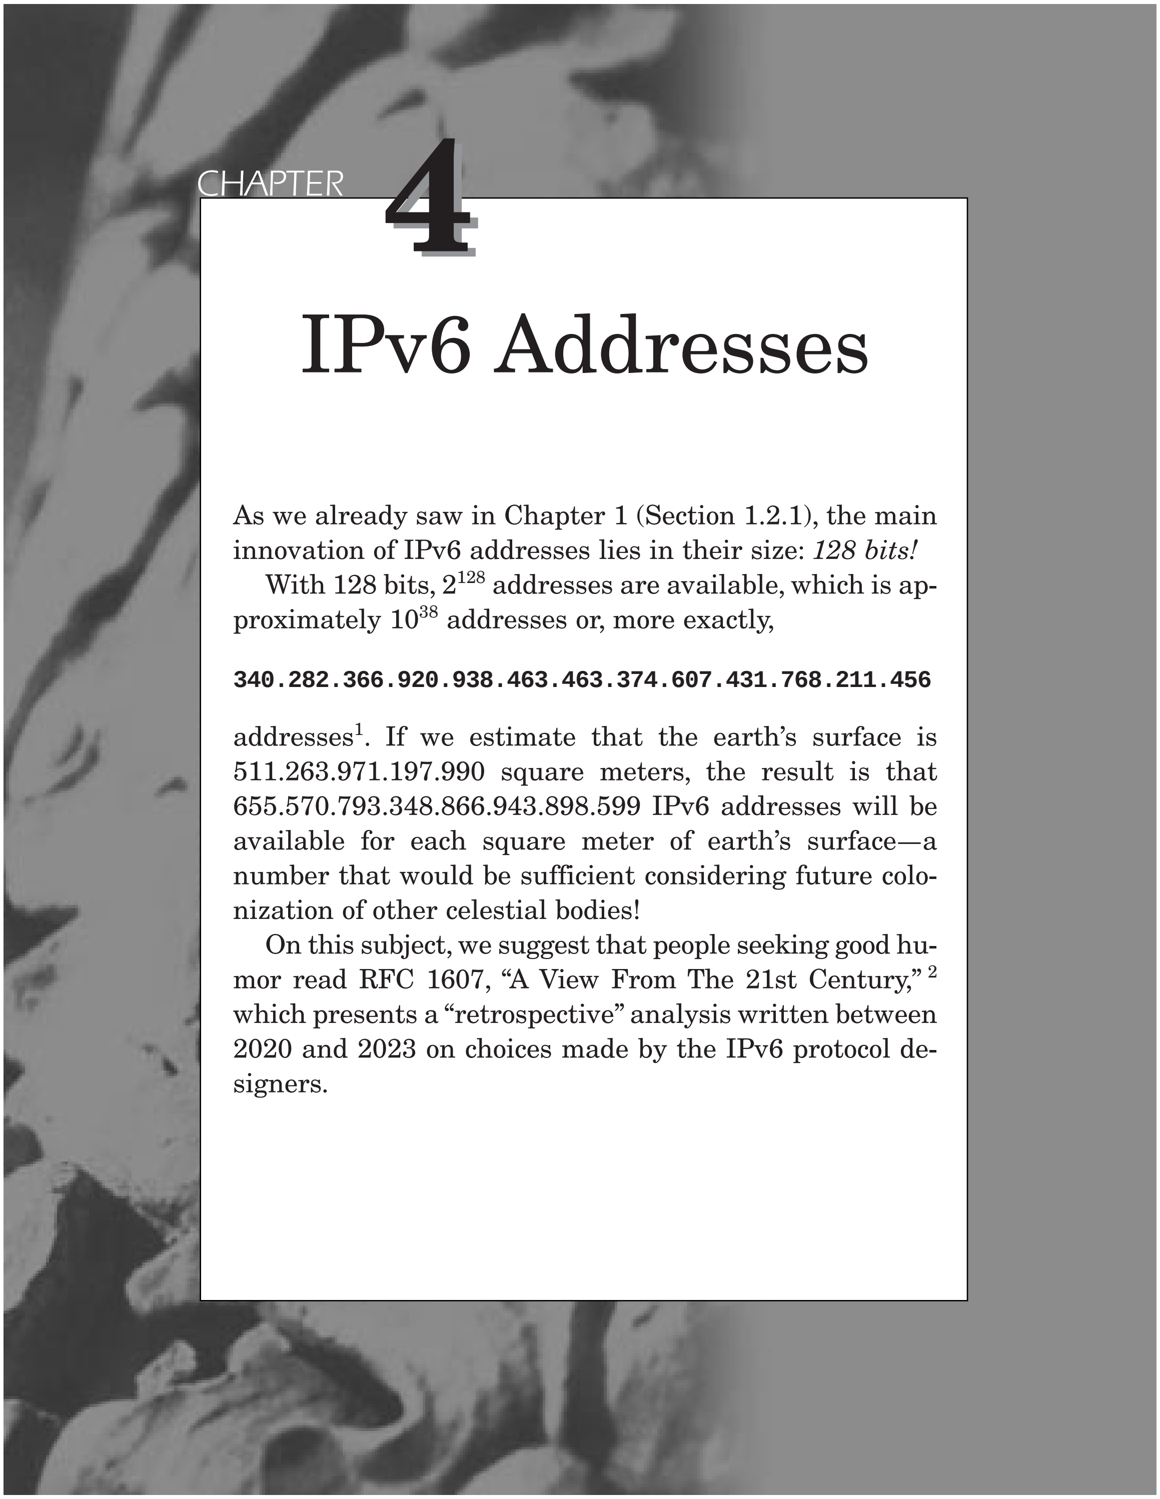

CHAPTER 4

IPv6 Addresses

As we already saw in Chapter 1 (Section 1.2.1), the main innovation of IPv6 addresses lies in their size: *128 bits!*

With 128 bits, 2^{128} addresses are available, which is approximately 10^{38} addresses or, more exactly,

340.282.366.920.938.463.463.374.607.431.768.211.456

addresses[1]. If we estimate that the earth's surface is 511.263.971.197.990 square meters, the result is that 655.570.793.348.866.943.898.599 IPv6 addresses will be available for each square meter of earth's surface—a number that would be sufficient considering future colonization of other celestial bodies!

On this subject, we suggest that people seeking good humor read RFC 1607, "A View From The 21st Century,"[2] which presents a "retrospective" analysis written between 2020 and 2023 on choices made by the IPv6 protocol designers.

4.1 The Addressing Space

IPv6 designers decided to subdivide the IPv6 addressing space on the basis of the value assumed by leading bits in the address; the variable-length field comprising these leading bits is called the *Format Prefix* (FP)[3]. The allocation scheme adopted is shown in Table 4-1.

Table 4-1

Allocation of the IPv6 addressing space

Allocation	Prefix (binary)	Fraction of Address Space
Reserved	0000 0000	1/256
Unassigned	0000 0001	1/256
Reserved for NSAP addresses	0000 001	1/128
Reserved for IPX addresses	0000 010	1/128
Unassigned	0000 011	1/128
Unassigned	0000 1	1/32
Unassigned	0001	1/16
Aggregatable global unicast addresses	001	1/8
Unassigned	010	1/8
Unassigned	011	1/8
Reserved for Geographic-based addresses	100	1/8
Unassigned	101	1/8
Unassigned	110	1/8
Unassigned	1110	1/16
Unassigned	1111 0	1/32
Unassigned	1111 10	1/64
Unassigned	1111 110	1/128
Unassigned	1111 1110 0	1/512
Link Local addresses	1111 1110 10	1/1024

From the first examination of the table, we can see that only 15 percent of the addressing space is initially used by IPv6, thus leaving 85 percent of the addressing space unassigned for future uses.

The format prefixes 001 through 111, except for Multicast Addresses (1111 1111), are all required to have 64-bit interface identifiers in EUI-64 format (see Section 4.10 for definitions).

Reserved addresses must not be confused with Unassigned addresses. They represent 1/256 of the addressing space (FP = 0000 0000) and are used for *unspecified* addresses (see Section 4.6.6), *loopback* (see Section 4.6.7), and *IPv6 with embedded IPv4* addresses (see Section 4.6.8).

Other reserved addresses are *NSAP* addresses (FP = 0000 001) that represent 1/128 of the addressing space and can be derived from ISO/OSI *Network Service Access Point* (NSAP) addresses. A proposal in this direction is specified by RFC 1888 [3] and described in Section 4.6.9.

In the same way, a space for *IPX* addresses is reserved (FP = 0000 010) equal to 1/128 of the addressing space. These addresses can be derived from Novell IPX addresses (see Section 4.6.10).

The last type of reserved address is the *Geographic-based* address (FP = 100), which is the most similar to the present IPv4 addresses from the management point of view. The Geographic-based address was conceived to be assigned to the end user on the basis of the user's geographic location. This kind of address didn't gain much popularity because it potentially causes the routing table's explosion problems mentioned in Section 1.2.6. Of the addressing space, 1/8 is reserved for Geographic-based addresses (see Section 4.6.3), but they have been removed from the last IETF draft on Addressing Architecture.

The following unicast addresses are certain to be used from the beginning:

- Aggregatable Global Unicast addresses (FP = 001), which represent 1/8 of the addressing space; they will be described in Section 4.6.2.
- Link Local addresses (FP = 1111 1110 10), which represent 1/1024 of the addressing space; they will be described in Section 4.6.4.
- Site Local addresses (FP = 1111 1110 11), which represent 1/1024 of the addressing space; they will be described in Section 4.6.5.
- Multicast addresses (FP = 1111 1111), which represent 1/256 of the addressing space; they will be described in Section 4.8.

4.2 Syntax of IPv6 Addresses

IPv4 addresses are 32 bits (4 octets) long. When they are written, each octet is the representation of an unsigned integer, and the 4 octets are written as four decimal numbers divided by three dots (. . .). For example:

```
130.192.1.143
```

For IPv6 addresses, defining a similar syntax is necessary, taking into account that IPv6 addresses are four times longer. The syntax standardized by RFC 1884 [3] recommends considering 128 bits (16 octets) of the IPv6 address as eight unsigned integers on 16 bits and writing each number with four hexadecimal digits; we divide each number from the preceding one and from the following one by using a colon (:). For example:

```
FEDC:BA98:7654:3210:FEDC:BA98:7654:3210
```

The preceding example clarifies the difficulty of the manual management of IPv6 addresses and the need for DHCP and DNS servers (as discussed in Section 2.13). Some IPv6 designers see some advantages in the users' difficulty remembering and writing IPv6 addresses: this way, users will be forced to use names more and more, and addresses will become a problem more internal to the network and functional to the routing of packets.

Nevertheless, the preceding example is not completely realistic; the following are more realistic examples of addresses:

```
1080:0000:0000:0000:0008:0800:200C:417A
0000:0000:0000:0000:0000:0000:0A00:0001
```

Clearly, more compressed forms of representation are easier for these kinds of addresses. One shortcut derives from the fact that we do not need to write the leading zeros in each group of digits; for example, we can write 0 instead of 0000, 1 instead of 0001, 20 instead of 0020, and 300 instead of 0300. If we apply this shortcut, the two preceding addresses become

```
1080:0:0:0:8:800:200C:417A
0:0:0:0:0:0:A00:1
```

A further simplification is represented by the symbol `::`, which replaces a series of zeros. By applying this shortcut, the two preceding addresses become

```
1080::8:800:200C:417A
::A00:1
```

Note that the preceding shortcut can be applied only once to an address. We make the assumption that the IPv6 address has a fixed length so that we can compute how many zeros have been omitted. This shortcut can be applied either to the center of the address (as in the case of the first address), or to the leading (as in the case of the second address) or trailing zeros.

If we consider the case of multicast, loopback, or unspecified addresses, we realize how useful this shortcut is. In fact, the extended form of these addresses results in the following:

FF01:0:0:0:0:0:0:43	A multicast address
0:0:0:0:0:0:0:1	The loopback address
0:0:0:0:0:0:0:0	The unspecified address

They can be represented in compressed form as follows:

FF01::43	A multicast address
::1	The loopback address
::	The unspecified address

A special case is valid for addresses such as **0:0:0:0:0:0:A00:**. The six leading zeros indicate that it is an IPv6 address with an embedded IPv4 address (see Section 4.6.8). In particular, this IPv6 address is associated with the IPv4 address **10.0.0.1**. Only in this case can a mixed IPv4/IPv6 notation be used. In its extended form, the resulting address is

```
0:0:0:0:0:0:10.0.0.1
```

and in compressed form, the address is

```
:::10.0.0.1
```

The representation of IPv6 prefixes is similar to the way IPv4 addresses' prefixes are written in CIDR notation. An IPv6 address prefix is represented by the notation

```
ipv6-address/prefix-length
```

where **ipv6-address** is any of the notations described in this section and **prefix-length** is a decimal value specifying the length of the prefix in bits.

For example, to indicate a subnet with an 80-bit prefix, we use the following notation:

```
1080:0:0:0:8::/80
```

Note that in this case the three central zeros cannot be eliminated because the notation :: has already been used once at the end of the address.

For example, the 60-bit prefix

```
12AB00000000CD3
```

has the following legal representations:

```
12AB:0000:0000:CD30:0000:0000:0000:0000/60

12AB::CD30:0:0:0:0/60
12AB:0:0:CD30::/60
```

However, the following representations are not legal:

12AB:0:0:CD3/60 Because we can drop leading zeros but not trailing zeros within any 16-bit chunk of the address

12AB::CD30/60 Because the address to the left of / expands to **12AB:0000:0000:0000:0000:000:0000:CD30**

12AB::CD3/60 Because the address to the left of / expands to **12AB:0000:0000:0000:0000:000:0000:0CD3**

The node address and its prefix can be combined as shown here. The lines

```
node address: 12AB:0:0:CD30:123:4567:89AB:CDEF
prefix: 12AB:0:0:CD30::/60
```

can be abbreviated as

```
12AB:0:0:CD30:123:4567:89AB:CDEF/60
```

4.3 Types of IPv6 Addresses

As we already saw in Section 2.2, interfaces are addressable in IPv6. More precisely, we can say that a 128-bit IPv6 address is associated with an interface or to a set of interfaces. In particular, RFC 1884 [3] identifies three types of IPv6 addresses:

- *Unicast:* This type is the address of a single interface. A packet forwarded to a unicast address is delivered only to the interface identified by that address.

- *Anycast:* This type is the address of a set of interfaces typically belonging to different nodes. A packet forwarded to an anycast address is delivered to only one interface of the set (the nearest to the source node, according to the routing metric).
- *Multicast:* This type is the address of a set of interfaces that typically belong to different nodes. A packet forwarded to a multicast address is delivered to all interfaces belonging to the set.

The main differences between IPv4 and IPv6 addresses are the appearance of anycast addresses in IPv6 and the disappearance of IPv4 broadcast addresses, replaced by IPv6 multicast addresses.

4.4 The Addressing Model

We have already learned that addresses belong to interfaces, not to nodes. A node can be identified by any unicast address associated with its interfaces. An IPv6 unicast address refers to a single interface. A single interface can be assigned more addresses of the same type or of different types (unicast, anycast, or multicast). The following are two exceptions to this model:

1. A single IPv6 address can be assigned to a group of interfaces belonging to a node if IPv6 implementation treats that group as a single interface when presenting packets to the IP layer. This capability is useful in fault tolerant systems, in which the presence of only one interface can represent a single point of failure, or to implement a mechanism for load sharing over multiple physical interfaces.
2. Routers can have *unnumbered* interfaces—that is, without any addresses. This can be the case for interfaces on point-to-point links where the presence of addresses is not essential. This setup can simplify a router's configuration, but its use is discouraged from the management point of view because explicitly referring to an interface is not possible if the interface is not associated with a unicast address.

IPv6 assumes that a subnet (or subnetwork, see Section 2.4) is associated with a link (or communication channel, see Section 2.2). More subnets can be associated with the same link, but a subnet cannot be associated with more than one link.

4.5 Assignment of IPv6 Addresses

We have already seen that IPv6 addresses will be unique at a worldwide level, and this uniqueness implies the existence of one or more organizations to assign these addresses.

RFC 1881 [6] specifies that the IPv6 addressing space must be managed in the Internet community's interest through a small central authority availing itself of the cooperation of peripheral authorities.

The Internet community decided that the appropriate entity to perform the role of central authority would be the *Internet Assigned Numbers Authority* (IANA). The IANA will base the IPv6 addressing space management on suggestions coming from the *Internet Architecture Board* (IAB) and from the *Internet Engineering Steering Group* (IESG).

The IANA will delegate to regional and other local registries the task of making specific address allocations to network service providers and other subregional registries. Individuals and organizations can obtain address allocations directly from the appropriate regional (or other) registry or from their service providers.

The IANA will try to prevent monopolies and instances of abuse.

The IANA will develop a plan for the initial IPv6 address allocation, including a provision for the automatic allocation of IPv6 addresses to holders of IPv4 addresses. IANA will also develop mediation and appeals procedures concerning delegation and revocation.

The IANA has already identified three local authorities to collaborate with for IPv6 address allocation:

- RIPE-NCC (Réseaux IP Européens Network Coordination Centre) for Europe
- INTERNIC (Internet Network Information Center) for Northern America
- APNIC (Asian and Pacific Network Information Center) for Asian and Pacific countries

4.6 Unicast Addresses

IPv6 unicast addresses are continuous, bit-wise, maskable addresses similar to IPv4 addresses with *Classless Inter-Domain Routing* (CIDR) [7], as described in Section 1.2.1. We have already seen that the following types of unicast addresses have been specified: Aggregatable Global Unicast, Geographic-based, IPv4, NSAP, IPX, Link Local, Site Local, nonspecified, and loopback. They will be described in this section. Additional address types will be defined later.

IPv6 nodes may have little knowledge of the internal structure of an IPv6 address. In the simplest case, a node may consider an IPv6 address as a 128-bit string (see Figure 4-1).

A slightly more sophisticated node may have a vision of the IPv6 address structured into two parts by means of the prefix that identifies the subnet (see Figure 4-2).

Routers can have even more sophisticated visions of the address and know other boundaries. The sophistication level of routers depends on what position routers hold in the routing hierarchy.

4.6.1 Example of a Unicast Address

An example of a unicast address format that will likely be common on LANs is the one that allows us to identify the node within the subnet from its 48-bit *MAC address*. Even if, until now, MAC addresses have been assigned on 48 bits, the *EUI-64* standard introduces MAC addresses on 64 bits to be used in the future (see Section 4.10). To be compliant with this standard, IPv6 uses identifiers on 64 bits from the beginning interface (see Figure 4-3).

The *subscriber ID* identifies the set of addresses allocated to a given organization. The *subnet ID* divides this set into several subnets (in this case, the prefix will be 64 bits). The 48-bit MAC address is extended to 64 bits using the EUI-64 rules, and the address is used to identify the inter-

Figure 4-1
IPv6 address non-structured vision

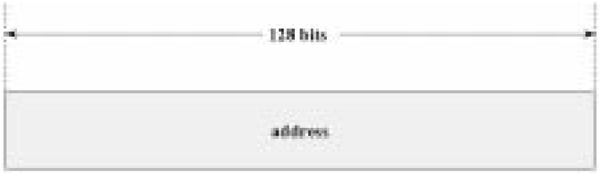

Figure 4-2
IPv6 address and prefix

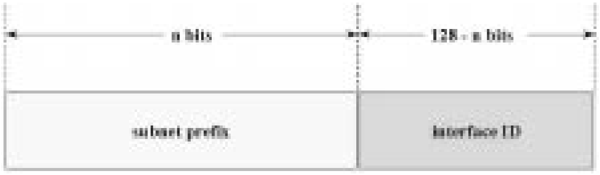

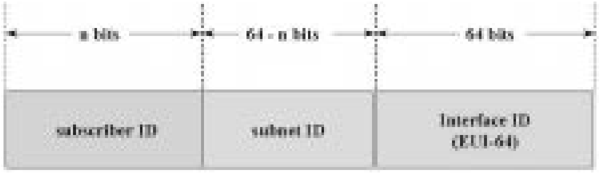

Figure 4-3 Example of a unicast address

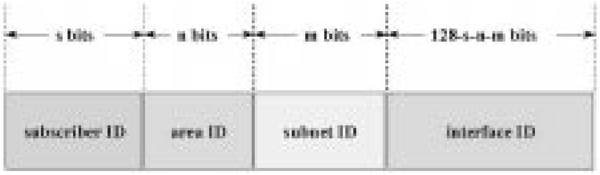

Figure 4-4 Two hierarchical levels

face within the subnet. The use of the MAC address makes possible a very simple form of address autoconfiguration: The interface can learn the first 64 bits from the router and autoconfigure its address by linking the 64 bits derived from its MAC address to them. In case the interface doesn't have a MAC address, other forms of layer 2 addresses can be used—for example, *E.164* addresses (ISDN numbers) for public network interfaces.

If the organization is particularly wide, it can decide that only one level of internal hierarchy is not enough and to configure two hierarchy levels: area and subnet. This solution is shown in Figure 4-4. Using an interface ID smaller than 64 bits may be desirable to leave more space for *area ID* and *subnet ID* fields.

Anyhow, the main partition remains the one between the interface address and the remaining part of the address. In fact, as we saw in Section 2.5, when a node forwards a packet, it checks whether the destination address is reachable through one of its interfaces—that is, whether the destination node is connected to one of its links. To execute this operation, knowing the length of the subnet prefix independently from existing hierarchical levels is essential. This number is

```
n = 128 - length(interface address)
```

according to the description in Figure 4-2.

4.6.2 Aggregatable Global Unicast Addresses

Aggregatable Global Unicast addresses are specified in *IP Version 6 Addressing Architecture* [16]. These addresses, which are characterized by FP = 001, are illustrated in Figure 4-5.

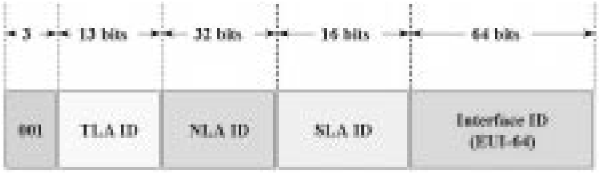

Figure 4-5
An Aggregatable Global Unicast address

The *Top-Level Aggregation IDentifier* (TLA ID) field is assigned to an organization providing public transit topology. It is specifically not assigned to an organization providing only leaf or private transit topology. The IANA will assign small blocks of TLA ID to IPv6 registries. At present, four registries exist; see Table 4-2.

The *Next-Level Aggregation IDentifier* (NLA ID) field is used by organizations assigned a TLA ID to create an addressing hierarchy and to identify sites (the ISP users).

The *Site-Level Aggregation IDentifier* (SLA ID) field is used by users assigned a TLA ID to create an addressing hierarchy within the sites, and this usually includes the subnet identifier.

This kind of assignment satisfies most users who can have at their disposal 64 thousand subnets, each one of practically unlimited size.

A discussion of problems related to Aggregatable Global Unicast addresses can be found in Section 7.6 and in RFC 1887 [8], where the connection between routing and addressing is examined, either within a domain or between different domains.

The Unicast addresses to be used in the IPv6 testing phase are detailed in Section A.4 of Appendix A.

4.6.3 Geographic-Based Addresses

Geographic-based addresses have been studied and proposed in the SIP project (see Section 1.5.4), but a final decision about them has not yet been made because ISPs strongly oppose them. In the latest IETF drafts, they are no longer present.

So that we can deploy these addresses, the world must be subdivided into continents, then into regions, and then into metropolitan areas. All ISPs that serve a given area must interconnect to route packets correctly. In this way, addresses can be directly allocated to end users who maintain the addresses even if they change ISPs. The ISPs' opposition is based on the complexity of routing table management.

Table 4-2

Current registries

Scope	Authority
Multiregional	IANA
Europe	RIPE-NCC
Northern America	INTERNIC
Asia and Pacific	APNIC

Geographic-based addresses have not yet been definitively abandoned, as shown by the fact that they have been allocated 1/8 of the IPv6 addressing space (FP = 100). Nevertheless, at the moment, there are no plans to use them.

For a discussion of advantages and drawbacks of Aggregatable Global Unicast and Geographic-based solutions, see Chapter 7.

4.6.4 Link Local Addresses

Link Local addresses (FP = 1111 1110 10) are designed to be used on each link for address autoconfiguration and for neighbor discovery functions. Their format is illustrated in Figures 4-6 and 4-7.

Suppose we have a small LAN with a few PCs connected and without a router; in this case, Link Local addresses turn out to be the only addresses we need.

Let's consider, for example, a PC with an IEEE 802.3 board with MAC address **08-00-02-12-34-56**. If we assume that the 48-bit MAC address is used as the interface ID, the PC's IPv6 Link Local address is

```
FE80:0000:0000:0000:0000:0800:0212:3456
```

or its compressed form is

```
FE80::800:0212:3456
```

In contrast, if we assume that the 64-bit EUI-64 (see Section 4.10) address is used as the interface ID, the PC's IPv6 Link Local address is

```
FE80:0000:0000:0000:0A00:02FF:FE12:3456
```

or its compressed form is

```
FE80::A00:2FF:FE12:3456
```

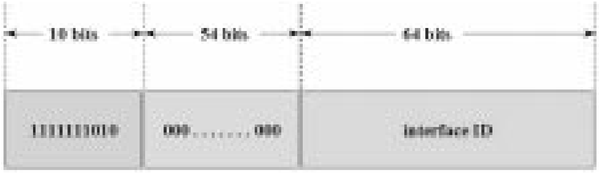

Figure 4-6 *Link Local addresses*

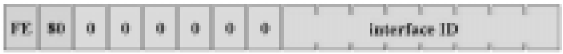

Figure 4-7 *Typical example of a Link Local address*

Remember that routers must never retransmit IPv6 packets that have a Link Local address as a source address.

4.6.5 Site Local Addresses

Site Local (FP = 1111 1110 11) addresses are designed to replace IPv4 addresses defined by RFCs 1597 [9] and 1918 [10] (see Section 1.3.3) for use in Intranets. Site Local addresses are therefore ideal for organizations not (yet) connected to the global Internet. They do not need any form of registration, and they have a format (see Figure 4-8) that makes replacing them with Aggregatable Global Unicast addresses simple when global connectivity to the Internet is desired.

The typical format of a Site Local address is illustrated in Figure 4-9.

A network using Site Local addresses can be complex because the presence of the subnet field on two octets allows us to have up to 64 thousand different subnets, each one with practically unlimited size.

A router with an IEEE 802.3 interface and MAC address **00-00-0C-12-34-56** connected to subnet 17 will have, on that interface, the following Site Local IPv6 address (using the 48-bit MAC address as the interface identifier):

```
FEC0:0000:0000:0000:0011:0000:0C12:3456
```

Its compressed form is as follows:

```
FEC0::11:0:C12:3456
```

If the EUI-64 MAC address is used (see Section 4.10) as the interface identifier, the resulting Site Local address is as follows:

```
FEC0:0000:0000:0011:0200:0CFF:FE12:3456
```

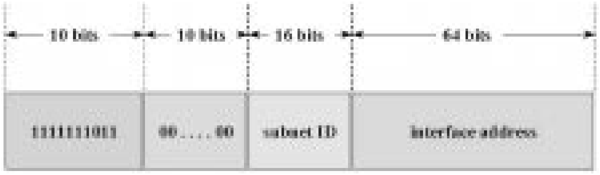

Figure 4-8
Site Local addresses

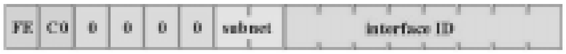

Figure 4-9
Typical example of a Site Local address

Here is its compressed form:

```
FEC0::11:200:CFF:FE12:3456
```

Again, remember that routers must never retransmit outside the site; IPv6 packets having a Site Local address as the source address. They must obviously retransmit these packets between different subnets of the same site.

4.6.6 The Unspecified Address

The address **`0000:0000:0000:0000:0000:0000:0000:0000`** is also called the *unspecified address,* and it can be written in the compressed form **`::`**. It must never be assigned to any interface because it indicates the absence of an IPv6 address. It can be used as a source address by a node during the configuration phase, when the node itself is trying to discover its IPv6 address. Also, the unspecified address must never be used as the destination address or in the Routing header (see Section 3.2.5).

4.6.7 The Loopback Address

The address **`0000:0000:0000:0000:0000:0000:0000:0001`** is also called the *loopback address* (its compressed form is **`::1`**), and it is used by a node to send an IPv6 packet to itself. It must never be assigned to any interface.

A node must never transmit outside itself any IPv6 packets with a loopback address as the source or destination address.

4.6.8 IPv6 Addresses with Embedded IPv4 Addresses

The transition mechanism from IPv4 to IPv6 includes a mechanism to dynamically tunnel IPv6 packets over the IPv4 routing infrastructure. (See Chapter 12 for details about the transition to IPv6.) IPv6 nodes that use this technique are assigned special IPv6 unicast addresses that carry an IPv4 address in the low-order 32 bits, as shown in Figure 4-10. These addresses are called *IPv4-compatible IPv6 addresses*.

An example of this type of address is the following:

```
::130.192.252.27
```

A second type of IPv6 address that holds an embedded IPv4 address is called an *IPv4-mapped IPv6 address* (see Figure 4-11). This second type of address is used to represent the address of an IPv4-only node in IPv6. An example of this type of address is the following:

```
::FFFF:130.192.252.27
```

4.6.9 NSAP Addresses

Today, the use of IPv6 addresses derived from ISO/OSI NSAP (FP = 0000 001) addresses is still under consideration, and a proposal in this direction is specified by RFC 1888 [4]. NSAP addresses are binary strings up to

Figure 4-10
IPv4-compatible IPv6 address

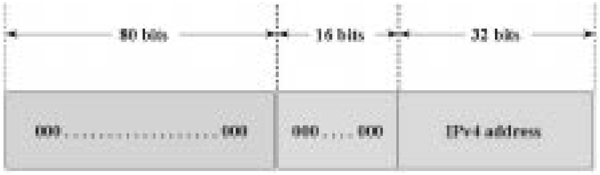

Figure 4-11
IPv4-mapped IPv6 address

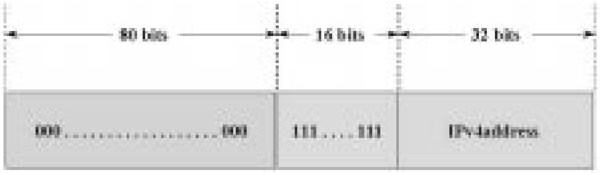

20 octets long defined in the OSI project by the standard ISO 8348 [12]. In the past, they held a certain interest because some organizations decided to adopt the layer 3 connectionless protocol ISO 8473 [11], which uses these addresses. NSAP addresses allow seven possible subformats, most of which are obsolete. Three subformats have been resumed and are used currently by the ATM [13] to address layer 2 ATM stations; they are illustrated in Figure 4-12.

At first glance, we can see that deriving IPv6 addresses starting from NSAP addresses (see Figure 4-13) clearly creates some problems because NSAP addresses (160 bits) are longer than IPv6 addresses (128 bits). These problems have three possible solutions:

1. To create a rule to map NSAP fields into IPv6 address fields; this solution is possible because not all NSAP fields have been used.
2. To truncate the NSAP and use it for routing while the complete NSAP address is transported inside a Destination option (see Section 3.2.8); for this purpose, a NSAPA option has been defined and is identified by the value 195 in the Option Type field (see Section 3.2.2).
3. To use a normal IPv6 address for the routing and to transport the complete NSAP inside a Destination option as in the previous case.

Considering the limited impact that, in our honest opinion, these types of addresses will have in the future, we will not discuss them further here. For a more detailed treatment, see RFC 1888 [4].

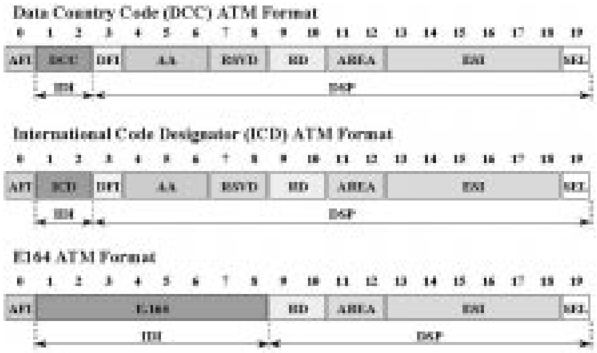

Figure 4-12
The three NSAP formats used by the ATM

Figure 4-13
IPv6 address drawn from a NSAP address

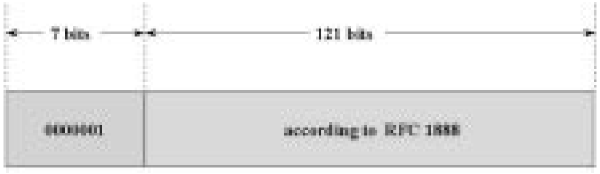

4.6.10 IPX Addresses

The network operating system Novell Netware is undoubtedly one of the most diffused in the field of PC networks. This network software can support several layer 3 (network) protocols, but the preferential protocol is *Internetwork Packet Exchange* (IPX) [14]. IPX is a connectionless protocol that assigns addresses to interfaces and is therefore very similar to IP. Addresses, which have the format shown in Figure 4-14, consist of two parts: Six octets contain the interface address (very frequently the MAC address), and four octets contain the segment ID.

The concept of *segments* is similar to the concept of *subnets* in IP. Because an IPX address is globally 80 bits long, implementing a relationship with IPv6 addresses (FP = 0000 010) that have 121 bits available for this purpose creates no problems (see Figure 4-15).

Nevertheless, at present no standard specifies how to implement this solution.

4.7 Anycast Addresses

We discussed the role of anycast addresses in Sections 1.3.2 and 4.3. Nevertheless, we must say that today we have little experience with the management of these addresses. Anycast addresses don't have separate addressing spaces (no particular FP value identifies anycast addresses); they simply are unicast addresses (belonging to one of the formats described in Section 4.6) assigned to more than one interface. When an anycast address is assigned to an interface, it must be explicitly configured to know that it is an anycast address; this information is usually specified by a qualifier at the time of the assignment.

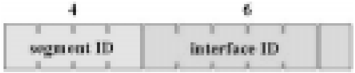

Figure 4-14
IPX address

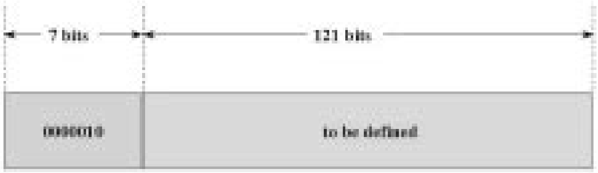

Figure 4-15
IPv6 address drawn from an IPX address

One possible use of anycast addresses is to identify a set of routers belonging to a given ISP, or all routers connected to a given subnet, or all border routers toward other domains.

For each anycast address, a prefix *P* identifies the topological region in which all interfaces belonging to that anycast address reside. Within this region—that is, this set of subnets—each interface associated with the anycast address must be advertised as a separate entry in a router's routing tables (see Section 2.6) so that the "nearest" interface belonging to the anycast set can be identified.

If the prefix P is null, the members of the set may have no topological locality. In this case, the anycast address must be advertised as a separate outing entry throughout the entire Internet, which presents a severe scaling limit on how many such "global" anycast sets can be supported.

After considering these "youthful" problems of anycast addresses, RFC 1884 [3] imposes the following two restrictions on the use of IPv6 anycast addresses:

- Anycast addresses must not be used as source addresses on IPv6 packets.
- Anycast addresses must not be assigned to IPv6 hosts—that is, they can be assigned to IPv6 routers only.

The only anycast address defined up till now is the *Subnet router anycast address;* its format is shown in Figure 4-16. Its intended use is to identify a set of routers connected to a given link. The *subnet prefix* must coincide with the prefix of the subnet associated with the link. A packet forwarded to the Subnet router anycast address will be delivered by a router connected to that link.

All routers are required to receive packets forwarded to the Subnet router anycast address on all the subnets on which they have interfaces.

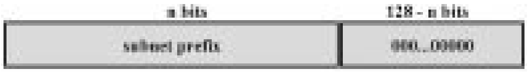

Figure 4-16
Anycast address

The Subnet router anycast address is useful, for example, either to solve the problem, present in IPv4, of the manual configuration of the default gateway on all hosts, or for a mobile host that needs to communicate with one of the routers on its home network.

4.8 Multicast Addresses

The possibility of implementing multicast transmissions on the Internet was developed in 1988 with the advent of class D IPv4 addresses. This feature is used widely by new multimedia applications that frequently need to transmit from one node to many nodes.

For this purpose, IPv6 specifies an addressing space identified by FP = 1111 1111; this format is illustrated in Figure 4-17.

The *flg* (flag) field is 4 bits long, and its structure is shown in Figure 4-18.

The first 3 bits are reserved for future uses and must be set to zero. The T bit can assume two different values:

- T = 0 indicates a permanently assigned (well-known) multicast address, assigned by the global Internet numbering authority (IANA).
- T = 1 indicates a transient multicast address, not permanently assigned.

The 4-bit *scp* (scope) field is used to limit the scope of the multicast group. Possible values for this field are indicated in Table 4-3.

The 112-bit *group ID* field identifies the multicast group, either permanent or transient, within a given scope. This means, for example, that equal ID groups can be simultaneously used in different parts of the network without interference, if their scopes are separate.

The meaning of a permanently assigned multicast address is independent of the scope value. Let's consider, for example, the *Network Time Protocol* (NTP) [15] servers group, which is the permanent group 43 hexadecimal

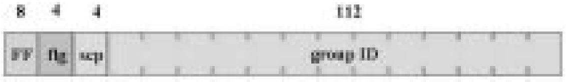

Figure 4-17
Multicast address

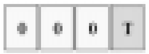

Figure 4-18
The flg field

Table 4-3

Allowed values for scp

scp	Meaning
0	Reserved
1	Node Local scope
2	Link Local scope
3	(Unassigned)
4	(Unassigned)
5	Site Local scope
6	(Unassigned)
7	(Unassigned)
8	Organization Local scope
9	(Unassigned)
A	(Unassigned)
B	(Unassigned)
C	(Unassigned)
D	(Unassigned)
E	Global scope
F	Reserved

assigned by IPv6. All the following four addresses belong to group 43, while having different meanings:

- FF01::43 means all NTP servers on the same node as the sender.
- FF02::43 means all NTP servers on the same link as the sender.
- FF05::43 means all NTP servers on the same site as the sender.
- FF0E::43 means all NTP servers present on the network.

Transient addresses can be associated with different applications in different parts of the network.

Moreover, multicast addresses must not be used as source addresses or appear in any Routing header (see Section 3.2.5).

4.8.1 Predefined Multicast Addresses

RFC 1884 [3] predefines a certain number of multicast addresses. They will be described in the following subsections.

4.8.1.1 RESERVED MULTICAST ADDRESSES The following multicast addresses are reserved for future uses:

```
FF00:0000:0000:0000:0000:0000:0000:0000
FF01:0000:0000:0000:0000:0000:0000:0000
FF02:0000:0000:0000:0000:0000:0000:0000
FF03:0000:0000:0000:0000:0000:0000:0000
FF04:0000:0000:0000:0000:0000:0000:0000
FF05:0000:0000:0000:0000:0000:0000:0000
FF06:0000:0000:0000:0000:0000:0000:0000
FF07:0000:0000:0000:0000:0000:0000:0000
FF08:0000:0000:0000:0000:0000:0000:0000
FF09:0000:0000:0000:0000:0000:0000:0000
FF0A:0000:0000:0000:0000:0000:0000:0000
FF0B:0000:0000:0000:0000:0000:0000:0000
FF0C:0000:0000:0000:0000:0000:0000:0000
FF0D:0000:0000:0000:0000:0000:0000:0000
FF0E:0000:0000:0000:0000:0000:0000:0000
FF0F:0000:0000:0000:0000:0000:0000:0000
```

4.8.1.2 ALL NODES ADDRESSES The following multicast addresses identify the group of all IPv6 nodes within the scope 1 (Node Local) and the scope 2 (Link Local):

```
FF01:0000:0000:0000:0000:0000:0000:0001
FF02:0000:0000:0000:0000:0000:0000:0001
```

4.8.1.3 ALL ROUTERS ADDRESSES The following multicast addresses identify the group of all IPv6 routers within the scope 1 (Node Local), the scope 2 (Link Local), and the scope 5 (Site Local):

```
FF01:0000:0000:0000:0000:0000:0000:0002
```

```
FF02:0000:0000:0000:0000:0000:0000:0002
FF05:0000:0000:0000:0000:0000:0000:0002
```

4.8.1.4 SOLICITED NODE MULTICAST ADDRESS

Multicast addresses in the range from

```
FF02:0000:0000:0000:0000:0001:FF00:0000
```

to

```
FF02:0000:0000:0000:0000:0001:FFFF:FFFF
```

are reserved for the Neighbor Discovery protocol (see Chapter 6) within the link. They are formed by taking the low-order 32 bits of the address (unicast or anycast) and appending them to the following prefix:

```
FF02:0000:0000:0000:0000:0001
```

For example, the Aggregatable Global Unicast address

```
2037::01:800:200E:8C6C
```

is associated with the Neighbor Discovery address (*Solicited Node Multicast Address*)

```
FF02::1:FF0E:8C6C
```

4.8.1.5 OTHER MULTICAST ADDRESSES

Other multicast addresses currently defined are as follows:

FF02:0:0:0:0:0:0:4	DVMRP Routers
FF02:0:0:0:0:0:0:5	OSPFIGP
FF02:0:0:0:0:0:0:6	OSPFIGP Designated Routers
FF02:0:0:0:0:0:0:7	ST Routers
FF02:0:0:0:0:0:0:8	ST Hosts
FF02:0:0:0:0:0:0:9	RIP Routers
FF02:0:0:0:0:0:0:A	EIGRP Routers
FF02:0:0:0:0:0:0:B	Mobile-Agents
FF02:0:0:0:0:0:0:D	All PIM Routers
FF02:0:0:0:0:0:0:E	RSVP-Encapsulation
FF02:0:0:0:0:0:1:1	Link Name
FF02:0:0:0:0:0:1:2	All-dhcp-agents
FF05:0:0:0:0:0:1:3	All-dhcp-servers
FF05:0:0:0:0:0:1:4	All-dhcp-relays
FF05:0:0:0:0:0:1:1000	

to FF05:0:0:0:0:0:1:13FF	Service Location
FF0X:0:0:0:0:0:0:100	VMTP Managers Group
FF0X:0:0:0:0:0:0:101	Network Time Protocol (NTP)
FF0X:0:0:0:0:0:0:102	SGI-Dogfight
FF0X:0:0:0:0:0:0:103	Rwhod
FF0X:0:0:0:0:0:0:104	VNP
FF0X:0:0:0:0:0:0:105	Artificial Horizons
FF0X:0:0:0:0:0:0:106	NSS - Name Service Server
FF0X:0:0:0:0:0:0:107	AUDIONEWS - Audio News
FF0X:0:0:0:0:0:0:108	SUN NIS+ Information Service
FF0X:0:0:0:0:0:0:109	MTP Multicast Transport Protocol
FF0X:0:0:0:0:0:0:10A	IETF-1-LOW-AUDIO
FF0X:0:0:0:0:0:0:10B	IETF-1-AUDIO
FF0X:0:0:0:0:0:0:10C	IETF-1-VIDEO
FF0X:0:0:0:0:0:0:10D	IETF-2-LOW-AUDIO
FF0X:0:0:0:0:0:0:10E	IETF-2-AUDIO
FF0X:0:0:0:0:0:0:10F	IETF-2-VIDEO
FF0X:0:0:0:0:0:0:110	MUSIC-SERVICE
FF0X:0:0:0:0:0:0:111	SEANET-TELEMETRY
FF0X:0:0:0:0:0:0:112	SEANET-IMAGE
FF0X:0:0:0:0:0:0:113	MLOADD
FF0X:0:0:0:0:0:0:114	any private experiment
FF0X:0:0:0:0:0:0:115	DVMRP on MOSPF
FF0X:0:0:0:0:0:0:116	SVRLOC
FF0X:0:0:0:0:0:0:117	XINGTV
FF0X:0:0:0:0:0:0:118	microsoft-ds
FF0X:0:0:0:0:0:0:119	nbc-pro
FF0X:0:0:0:0:0:0:11A	nbc-pfn
FF0X:0:0:0:0:0:0:11B	lmsc-calren-1
FF0X:0:0:0:0:0:0:11C	lmsc-calren-2
FF0X:0:0:0:0:0:0:11D	lmsc-calren-3
FF0X:0:0:0:0:0:0:11E	lmsc-calren-4
FF0X:0:0:0:0:0:0:11F	ampr-info
FF0X:0:0:0:0:0:0:120	mtrace
FF0X:0:0:0:0:0:0:121	RSVP-encap-1
FF0X:0:0:0:0:0:0:122	RSVP-encap-2
FF0X:0:0:0:0:0:0:123	SVRLOC-DA
FF0X:0:0:0:0:0:0:124	rln-server
FF0X:0:0:0:0:0:0:125	proshare-mc
FF0X:0:0:0:0:0:0:126	dantz
FF0X:0:0:0:0:0:0:127	cisco-rp-announce
FF0X:0:0:0:0:0:0:128	cisco-rp-discovery
FF0X:0:0:0:0:0:0:129	gatekeeper
FF0X:0:0:0:0:0:0:12A	iberiagames
FF0X:0:0:0:0:0:0:202	SUN RPC PMAPPROC_CALLIT
FF0X:0:0:0:0:0:2:0000 to FF0X:0:0:0:0:0:2:7FFD	Multimedia Conference Calls
FF0X:0:0:0:0:0:2:7FFE	SAPv1 Announcements
FF0X:0:0:0:0:0:2:8000 to FF0X:0:0:0:0:0:2:FFFF	SAP Dynamic Assignments

4.9 Which Addresses for a Node?

A reasonable question at this point is: Which addresses must a node have? The answer comes, once again, from RFC 1884, which lists all addresses that an IPv6 node can have.

4.9.1 Addresses of a Host

A host is required to recognize the following addresses as identifying itself:

- Its Link Local address for each interface
- Unicast addresses assigned to interfaces
- The loopback address
- All-Nodes multicast address
- Neighbor Discovery multicast addresses associated with all unicast and anycast addresses assigned to interfaces
- Multicast Addresses of groups to which the node belongs

4.9.2 Addresses of a Router

A router is required to recognize the following addresses as identifying itself:

- Its Link Local address for each interface
- Unicast addresses assigned to interfaces
- The loopback address
- The Subnet Router anycast address for all links on which it has interfaces
- Other anycast addresses assigned to interfaces
- All-nodes multicast address
- All-routers multicast address
- Neighbor Discovery multicast addresses associated with all unicast and anycast addresses assigned to interfaces
- Multicast addresses of groups to which the node belongs

4.10 The EUI-64 Interface Identifier

The IEEE has introduced a new type of MAC address, 64-bits long, called the EUI-64.

Until now, MAC addresses have been on 48 bits: 24 bits assigned by the IEEE and 24 bits manufacturer selected. The 24 bits assigned by the IEEE are called *Organization Unique Identifier* (OUI). Any company that has received an OUI from the IEEE can use it also for the new EUI-64 identifiers. It is sufficient to use the OUI as the first 24 bits and append them to the 40 manufacturer-selected bits.

Mapping the old 48-bit MAC addresses to a new 64-bit representation is also possible. The mapping process consists of inserting two octets with the value 0xFF and 0xFE between the OUI and the manufacturer-selected bits.

To obtain an IPv6 interface identifier from an EUI-64 address, we must complement the Universal/Local bit—that is, the next-to-last bit of the first octet.

The mapping of Universal MAC addresses to IPv6 interface identifiers is illustrated in Figure 4-19 for 48-bit MAC addresses and in Figure 4-20 for EUI-64.

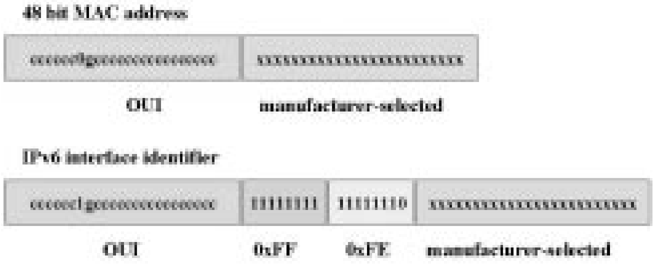

Figure 4-19
Address mapping from 48-bit to IPv6

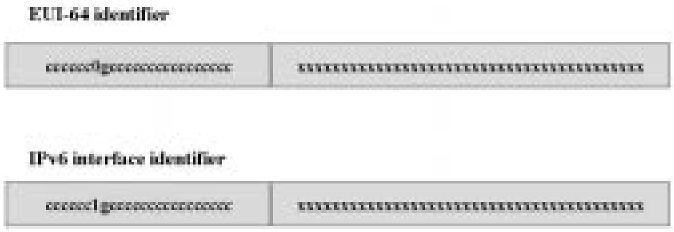
EUI-64 identifier
cccccc0gcccccccccccccccc
xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
IPv6 interface identifier
cccccc1gcccccccccccccccc
xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx

Figure 4-20
Address mapping from EUI to IPv6

REFERENCES

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

[2] V. Cerf, *RFC 1607: A View From The 21st Century*, April 1994.

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

[4] J. Bound, B. Carpenter, D. Harrington, J. Houldsworth, A. Lloyd*, RFC 1888: OSI NSAPs and IPv6*, August 1996.

[5] C. Huitema, *IPv6: the new Internet Protocol*, Prentice-Hall, 1996.

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

[7] Y. Rekhter, T. Li, *RFC 1518: An Architecture for IP Address Allocation with CIDR*, September 1993.

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

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

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

[11] ISO/IEC 8473, *IS8473: Data communications protocol for providing the connectionless-mode network service*, 1988.

[12] ISO/IEC 8348, *IS8348: Annex A, Network Layer Addressing, and Annex B, Rationale for the material in Annex A*, 1993 (same as CCITT X.213, 1992).

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

[14] Matthew Naugle, *Network Protocol Handbook,* McGraw-Hill, 1994.

[15] D.L. Mills*, RFC 1305: Network Time Protocol (Version 3) Specification, Implementation*, March 1992.

[16] R. Hinden, S. Deering, *IP Version 6 Addressing Architecture,* Internet Draft, July 1997.