ADDRESS CAPABILITIES

IPv4 Addressing and Issues IPv4 addresses can be from an officially assigned public range or from an internal intranet private (but not globally unique) block. As noted, IPv4 theoretically allows up to 232 addresses, based on a four-octet address space. Hence, there are 4,294,967,296 unique values, which can be considered as a sequence of 256 "/8s," where each "/8" corresponds to 16,777,216 unique address values. Public, globally unique addresses are assigned by IANA. IP addresses are addresses of network nodes at layer 3; each device on a network (whether the Internet or an intranet) must have a unique address. In IPv4, it is a 32-bit (4-byte) binary address used to identify a host's network ID. It is represented by the nomenclature a.b.c.d (each of a, b, c, and d being from 1 to 255) (0 has a special meaning). Examples are 167.168.169.170, 232.233.229.209, and 200.100.200.100. The problem is that during the 1980s, many public, registered addresses were allocated to firms and organizations without any consistent control. As a result, some organizations have more addresses that they actually might need, giving rise to the present dearth of available "registerable" layer 3 addresses. Furthermore, not all IP addresses can be used due to the fragmentation described above. One approach to the issue would be a renumbering and a reallocation of the IPv4 addressing space. However, this is not as simple as it appears since it requires worldwide coordination efforts. Moreover, it would still be limited for the human population and the quantity of devices that will be connected to Internet in the medium-term future. At this juncture, and as a temporary and pragmatic approach to alleviate the dearth of addresses, network address translation (NAT) mechanisms are employed by organizations and even home users. This mechanism consists of using only a small set of public IPv4 addresses for an entire network to access the Internet. The myriad of internal devices are assigned IP addresses from a specifically designated range of Class A or Class C address that are locally unique but are duplicatively used and reused within various organizations. In some cases (e.g., residential Internet access use via Digital Subscriber Line [DSL] or cable), the legal IP address is only provided to a user on a time-lease basis, rather than permanently. Internal intranet addresses may be in the ranges 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16. In the internal intranet private address case, a NAT function is employed to map the internal addresses to an external public address when the private-to-public network boundary is crossed. This, however, imposes a number of limitations, particularly since the number of registered public addresses available to a company is almost invariably much smaller (as small as 1) than the number of internal devices requiring an address. A number of protocols cannot travel through a NAT device, and hence the use of NAT implies that many applications (e.g., VoIP) cannot be used effectively in all instances. As a consequence, these applications can only be used in intranets. Examples include: Multimedia applications such as videoconferencing, VoIP, or video-ondemand/IPTV do not work smoothly through NAT devices. Multimedia applications make use of real-time transport protocol (RTP) and real-time control protocol (RTCP). These in turn use User Datagram Protocol (UDP) with dynamic allocation of ports and NAT does not directly support this environment. IPsec is used extensively for data authentication, integrity, and confidentiality. However, when NAT is used, IPsec operation is impacted, since NAT changes the address in the IP header. Multicast, although possible in theory, requires complex configuration in a NAT environment and hence, in practice, is not utilized as often as could be the case. The need for obligatory use of NAT disappears with IPv6. 7.2.2 IPv6 Address Space The IPv6 addressing architecture is described in RFC 4291 February 2006 (12). One of the major modifications in the addressing scheme in IPv6 is a change to the basic types of addresses and how they are utilized. Unicast addresses are utilized for a majority of traditional (enterprise) communications, as was the case in IPv4. However, Broadcast as a specific addressing type has been eliminated; in its place support for multicast addressing has been expanded and made a required part of the protocol. A new type of addressing called anycast has also been implemented. In addition, there are a number of special IPv6 addresses. Figure 7.1 compares the two

address formats. Figure 7.2 provides a pictorial comparison of these three transmission (and address) modes. Logically, one can interpret the types of transmissions as follows3: Unicast transmission: “send to this one specific address” Multicast transmission: “send to every member of this specific group” Anycast transmission: “send to any one member of this specific group.” Typically (motivated by efficiency goals), the transmission occurs to the closest (in routing terms) member of the group. Generally one interprets anycast to mean “send to the closest member of this specific group.” ETSI standards on the M2M system require support for anycast, unicast, multicast and broadcast communication modes; whenever possible, a global broadcast is

Unicast • Brodacast • Multicast • Multicast • Anycast (new) • Special • Special • Unicast - An ID for an interface - Explict assignment - Limited: 255.255.255.255 - Directed: 11.1 - An ID For a set of interfaces. - Deliver to all of them class D: - An ID for a set of interfaces. Deliver to the nearest one. - Undistinguishable from unicast 224.0.0.0 - 239.255.255.255 - 0.0.0.0, 127.0.0.1 - ::, ::1 IPv4 IPv6 - N per interface - Based on IEEE EUI-64 FIGURE

expected to be replaced by a multicast or anycast in order to minimize the load on the communication network (13). The format of IPv6 addressing is described in RFC 2373. As noted, an IPv6 address consists of 128 bits, rather than 32 bits as with IPv4 addresses; the number of bits correlates to the address space, as follows: IP Version Size of Address Space IPv6 128 bits, which allows for 2128 or 340,282,366,920,938,463,463,374,607,431,768,211,456 (3.4 × 1038) possible addresses IPv4 32 bits, which allows for 232 or 4,294,967,296 possible addresses The relatively large size of the IPv6 address is designed to be subdivided into hierarchical routing domains that reflect the topology of the modern-day Internet. The use of 128 bits provides multiple levels of hierarchy and flexibility in designing hierarchical addressing and routing. The IPv4-based Internet currently lacks this flexibility (14). The IPv6 address is represented as eight groups of 16 bits each, separated by the “:” character. Each 16-bit group is represented by 4 hexadecimal digits, that is, each digit has a value between 0 and f (0,1, 2, . . . a, b, c, d, e, f with a = 10, b = 11, and so on, to f = 15). What follows is an IPv6 address example 3223:0ba0:01e0:d001:0000:0000:d0f0:0010 An abbreviated format exists to designate IPv6 addresses when all endings are 0. For example 3223:0ba0:: is the abbreviated form of the following address: 3223:0ba0:0000:0000:0000:0000:0000:0000 Similarly, only one 0 is written, removing 0’s in the left side, and four 0’s in the middle of the address. For example the address 3223:ba0:0:0:0:0::1234 is the abbreviated form of the following address 3223:0ba0:0000:0000:0000:0000:0000:1234

There is also a method to designate groups of IP addresses or subnetworks that is based on specifying the number of bits that designate the subnetwork, beginning from left to right, using remaining bits to designate single devices inside the network. For example, the notation 3223:0ba0:01a0::/48 indicates that the part of the IP address used to represent the subnetwork has 48 bits. Since each hexadecimal digit has 4 bits, this points out that the part used to represent the subnetwork is formed by 12 digits, that is: “3223:0ba0:01a0.” The remaining digits of the IP address would be used to represent nodes inside the network. As noted, anycast addresses are a new type of address defined in IPv6 (as originally defined in RFC 1546). The purpose of the anycast address functionality is to enable capabilities that were difficult to implement in IPv4 environments. Datagrams sent to the anycast address are automatically delivered to the device in the network that is the easiest to reach. Anycast addresses can be used to define a group of devices, any one of which can support a service request from the user sent to a single specific IP address. One example is situations where one needs a service that can be provided by a set of different (dispersed) servers, but where one does not specifically care which one provides it; a specific example here may be an Internet or video (streaming) cache. Another example of anycast addressing is a router arrangement that allows datagrams to be transmitted to whichever router in a group of equivalent routers is closest to the point of transmission; a specific example here may be to allow load sharing between routers. It should be noted that there is no special anycast addressing

format: anycast addresses are the same as unicast addresses from an address format perspective. In practicality, an anycast address is defined and created in a self-declarative manner when a unicast address is assigned to more than one device interface. Special IPv6 addresses, as follows (see Table 7.2 for additional details) (15): Auto-return or loopback virtual address. This address is specified in IPv4 as the 127.0.0.1 address. In IPv6, this address is represented as ::1. Not specified address (::). This address is not allocated to any node since it is used to indicate absence of address. IPv6 over IPv4 dynamic/automatic tunnel addresses. These addresses are designated as IPv4-compatible IPv6 addresses and allow the sending of IPv6 traffic over IPv4 networks in a transparent manner. They are represented as, for example, ::156.55.23.5. IPv4 over IPv6 addresses automatic representation. These addresses allow for IPv4-only nodes to still work in IPv6 networks. They are designated as “mapped from IPv4 to IPv6 addresses” and are represented as ::FFFF:, for example ::FFFF.156.55.43.3.