## IP Address Classes and CIDR

IP addresses are divided in to 'classes', based on the decimal value represented in the first octet. This class definition is referred to as the First Octet Rule. There are five classes of IP addresses: classes A, class B, class C, class D; and class E, but only class A, B and C addresses are used to identify devices connected to the Internet. Class D addresses are used for multicasting, and Class E addresses are reserved for experimental use. The subnet mask is related to the IP address class. Thus, once the IP address class is known, the default routing mask is also known. The IP address classes and their related subnet masks are: . Class A addresses range from 0.0.0.0 through 126.255.255.255 and use a default subnet mask of 255.0.0.0. In Class A addresses, the first octet is used as for the network ID while the last three octets are used for the host ID. In other words, the first 8 bits of the subnet mask are all 1s, hence a subnet mask of 255.0.0.0. As a result, networks that use Class A addresses can theoretically support a maximum of 256 networks and 16,581,375 (255x255x255) hosts, however, the first and the last address cannot be used. The first address is the network address and the last address is the broadcast address. For example, a network with an IP address of 10.10.11.12 has a network ID of 10.0.0.0, the fist address, and a broadcast address of 10.255.255.255, the last address. Thus networks with a Class A IP address space can support a maximum of 254 networks (28-2) and 16,777,214 hosts (224-2). Consequently, Class A addresses are used for a few networks with a very large number of hosts on each network. . Class B addresses range from 128.0.0.0 through 168.255.255.255 and 170.0.0.0 through 191.255.255.255. These addresses use a default subnet mask of 255.255.0.0. In Class B addresses, the first two bits are used as for the network ID while the last two bits are used for the host ID. As a result, networks that use Class B addresses can support a maximum of 65,534 networks (216-2) and 65,534 hosts. Consequently, Class B addresses are used for a reasonable number of medium sized networks. Note: IP addresses with a first octet of 127, i.e. 127.0.0.0 through 127.255.255.255 do not fall in either the Class A address range or the Class B address range. IP addresses that have a first octet of 127 are reserved for diagnostics purposes. . Class C addresses range from 192.0.0.1 through 223.225.225.225 and default subnet mask of 255.255.255.0. In Class C addresses, the first three bits are used as for the network ID while only the last bit is used for the host ID. As a result, networks that use Class C addresses can support a maximum of 16,777,214 networks and 254 hosts. Consequently, Class C addresses are used for a large number of networks with a relatively small number of hosts on each network. . Class D addresses are in the range 224.0.0.0 through 239.255.255.255. These addresses are reserved for multicast transmissions. . Class E addresses are in the range 240.0.0.0 through 254.255.255.255. These addresses are reserved for experimental use. Note: InterNIC has reserved a number of IP address range, including 169.0.0.1 through 169.253.255.254, which has been reserved by for future use; 0.0.0.0, which was originally defined for use as a broadcast address; and 127.0.0.0, which is used as the loopback address. 128.0.0.0, 191.255.0.0, 192.0.0.0, and 223.255.255.0 also are reserved. Classless Interdomain Routing (CIDR) Notation Class-based IP addressing is fairly rigid. Thus, a small company with 50 hosts that wants to connect to the Internet would need a Class C address. However, a Class C address range supports 253 hosts; therefore 203 addresses would be wasted. Similarly, a company with 4,000 hosts would require a Class B address to connect to the Internet. A Class B address can support up to 65,023 hosts, resulting in 61,023 addresses being wasted. This problem can be overcome by extending the default subnet mask by adding more continuous 1s to it. The result is that the network can support less hosts. Thus, the company that has 4,000 hosts would use a Class B address with a subnet mask of 255.255.240.0. This is achieved by extending the subnet mask by 4 bits so that the first 20 bits represent the network ID and 12 bits only represent the host ID. Thus the address range now supports only 4,094 hosts, representing a loss of only 94 addresses. We can calculate the number of hosts supported by using the formula: 2n-2 where n is the number of bits used for the host ID. We need to subtract 2 addresses: the network address and the broadcast address. In this example, 12 bits are used for the host ID. Thus using the formula we can see that this subnet mask supports 4,094 hosts (212-2). We can calculate the number of subnets supported by a subnet mask by using the same formula: 2n-2. However, this time n is the number of bits used for the network ID. We again need to subtract 2 addresses: the network address and the broadcast address. Thus, in the example 255.255.240.0, 20 bits represent the network ID therefore this subnet mask supports 1048,574 subnets (220-2). This solves the problem of IP address allocation on the internet but presents a problem for routing tables, as the routing table cannot determine the subnet mask on the basses of the IP address class. Hence a different format of representing the IP address and its subnet mask is required. This format is called the Classless Interdomain Routing (CIDR) notation, or prefix notation. CIDR is in essence an adaptation of the Dotted Decimal Format and represents the subnet mask as a number of bits used for the network ID. This number of bits is indicated after the IP address by the number that follows the slash (/) symbol. For example, the CIDR notation IP address 140.12.2.128/20 indicates that the first 20 bits of the IP address is used for the subnet mask, i.e., the first 20 bits are all 1s. Thus, the format as 255.255.240.0. In addition, the routing protocols must send the mask with the routing update. Variable-Length

Subnet Masks CIDR is used within the Internet. Its counterpart within an organization is the Variable-length subnet mask (VLSM). Like CIDR, VLSM allows you to allocate the required host bits on a granular basis. In other words, it allows you to provide only the bits required to address the number of hosts on a particular subnetwork. Like CIDR, VLSM requires a routing protocol that supports the sending of the subnet mask in its updates. The routing protocols that support VLSM are: RIPv2; OSPF; IS-IS; EIGRP; and BGP-4. The routing protocols do not support VLSM are: RIPv1; IGRP; and EGP.