TCP/IP Overview

Document ID:13769
Updated: Aug 10, 2005
Print

Contents

Introduction

In the two decades since their invention, the heterogeneity of networks has expanded further with the deployment of Ethernet, Token Ring, Fiber Distributed Data Interface (FDDI), X.25, Frame Relay, Switched Multimegabit Data Service (SMDS), Integrated Services Digital Network (ISDN), and most recently, Asynchronous Transfer Mode (ATM). The Internet protocols are the best proven approach to internetworking this diverse range of LAN and WAN technologies.

The Internet Protocol suite includes not only lower-level specifications, such as Transmission Control Protocol (TCP) and Internet Protocol (IP), but specifications for such common applications as electronic mail, terminal emulation, and file transfer. Figure 1 shows the TCP/IP protocol suite in relation to the OSI Reference model. Figure 2 shows some of the important Internet protocols and their relationship to the OSI Reference Model. For information on the OSI Reference model and the role of each layer, please refer to the document Internetworking Basics.

The Internet protocols are the most widely implemented multivendor protocol suite in use today. Support for at least part of the Internet Protocol suite is available from virtually every computer vendor.

TCP/IP Technology

This section describes technical aspects of TCP, IP, related protocols, and the environments in which these protocols operate. Because the primary focus of this document is routing (a layer 3 function), the discussion of TCP (a layer 4 protocol) will be relatively brief.

TCP

TCP is a connection-oriented transport protocol that sends data as an unstructured stream of bytes. By using sequence numbers and acknowledgment messages, TCP can provide a sending node with delivery information about packets transmitted to a destination node. Where data has been lost in transit from source to destination, TCP can retransmit the data until either a timeout condition is reached or until successful delivery has been achieved. TCP can also recognize duplicate messages and will discard them appropriately. If the sending computer is transmitting too fast for the receiving computer, TCP can employ flow control mechanisms to slow data transfer. TCP can also communicates delivery information to the upper-layer protocols and applications it supports. All these characteristics makes TCP an end-to-end reliable transport protocol. TCP is specified in RFC 793 .

Figure 1 – TCP/IP Protocol Suite in Relation to the OSI Reference Model

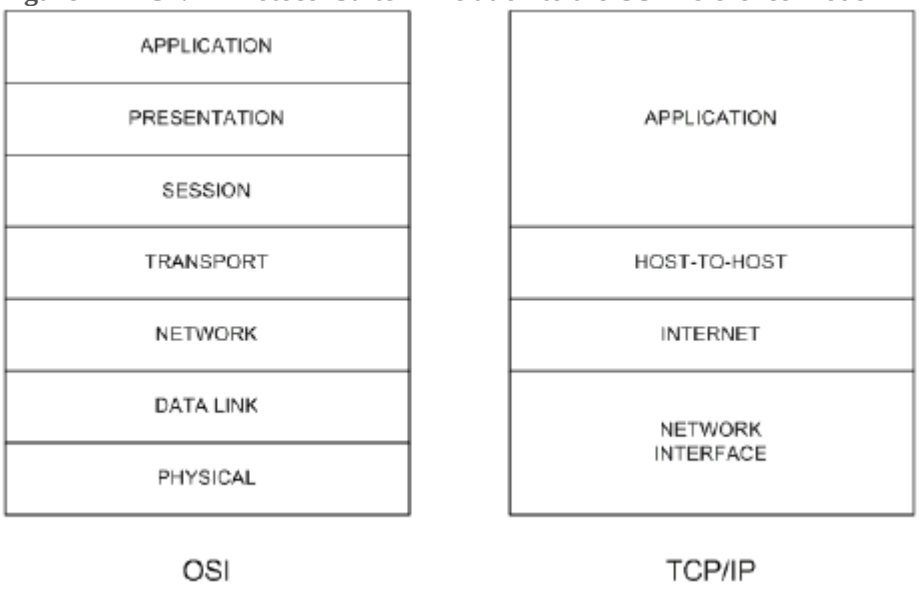

Figure 2 – Important Internet Protocols in Relation to the OSI Reference Model

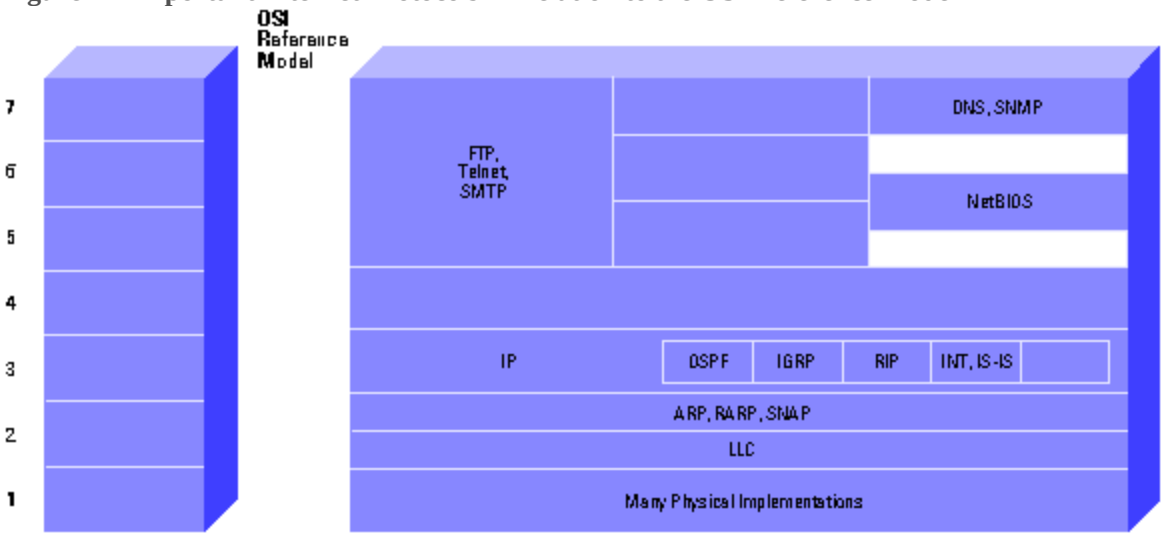

Refer to the TCP section of Internet Protocols for more information.

IP

IP is the primary Layer 3 protocol in the Internet suite. In addition to internetwork routing, IP provides error reporting and fragmentation and reassembly of information units called datagrams for transmission over networks with different maximum data unit sizes. IP represents the heart of the Internet Protocol suite.

Note: The term IP in the section refers to IPv4 unless otherwise stated explicitly.
IP addresses are globally unique, 32-bit numbers assigned by the Network Information Center. Globally unique addresses permit IP networks anywhere in the world to communicate with each other.

An IP address is divided into two parts. The first part designates the network address while the second part designates the host address.

The IP address space is divided into different network classes. Class A networks are intended mainly for use with a few very large networks, because they provide only 8 bits for the network address field. Class B networks allocate 16 bits, and Class C networks allocate 24 bits for the network address field. Class C networks only provide 8 bits for the host field, however, so the number of hosts per network may be a limiting factor. In all three cases, the left most bit(s) indicate the network class. IP addresses are written in dotted decimal format; for example, 34.0.0.1. Figure 3shows the address formats for Class A, B, and C IP networks.

Figure 3 – Address Formats for Class A, B, and C IP Networks

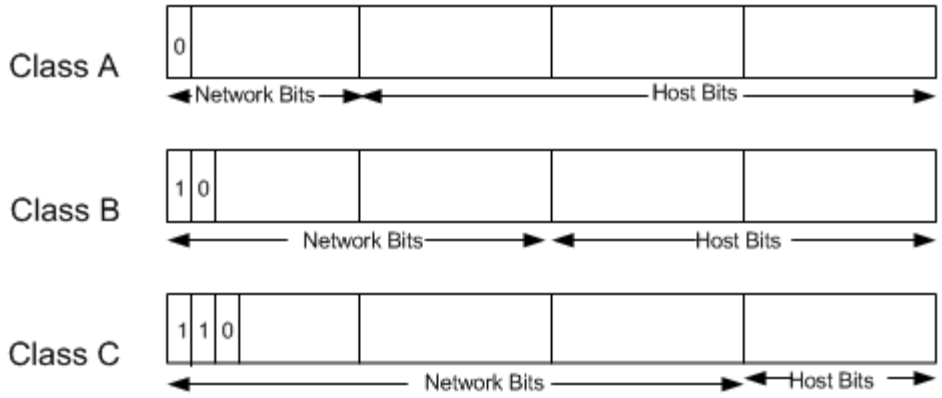

IP networks also can be divided into smaller units called subnetworks or "subnets." Subnets provide extra flexibility for the network administrator. For example, assume that a network has been assigned a Class A address and all the nodes on the network use a Class A address. Further assume that the dotted decimal representation of this network's address is 34.0.0.0. (All zeros in the host field of an address specify the entire network.) The administrator can subdivide the network using subnetting. This is done by "borrowing" bits from the host portion of the address and using them as a subnet field, as depicted in Figure 4.

Figure 4 – "Borrowing" Bits

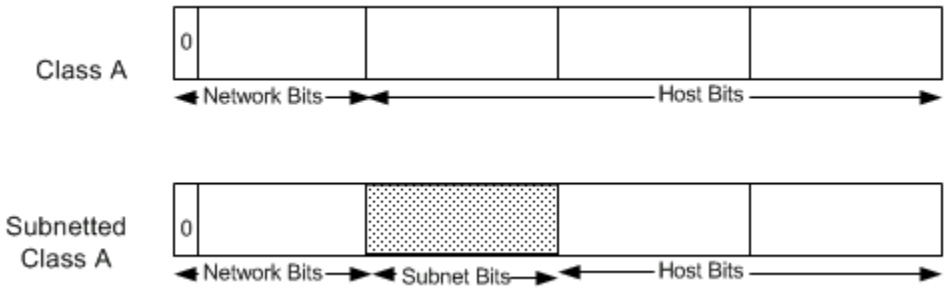

If the network administrator has chosen to use 8 bits of subnetting, the second octet of a Class A IP address provides the subnet number. In our example, address 34.1.0.0 refers to network 34, subnet 1; address 34.2.0.0 refers to network 34, subnet 2, and so on.

The number of bits that can be borrowed for the subnet address varies. To specify how many bits are used to represent the network and the subnet portion of the address, IP provides subnet masks. Subnet masks use the same format and representation technique as IP addresses. Subnet masks have ones in all bits except those that specify the host field. For example, the subnet mask that specifies 8 bits of subnetting for Class A address 34.0.0.0 is 255.255.0.0. The subnet mask that specifies 16 bits of subnetting for Class A address 34.0.0.0 is 255.255.255.0. Both of these subnet masks are pictured in Figure 5. Subnet masks can be passed through a network on demand so that new nodes can learn how many bits of subnetting are being used on their network.

Figure 5 – Subnet Masks

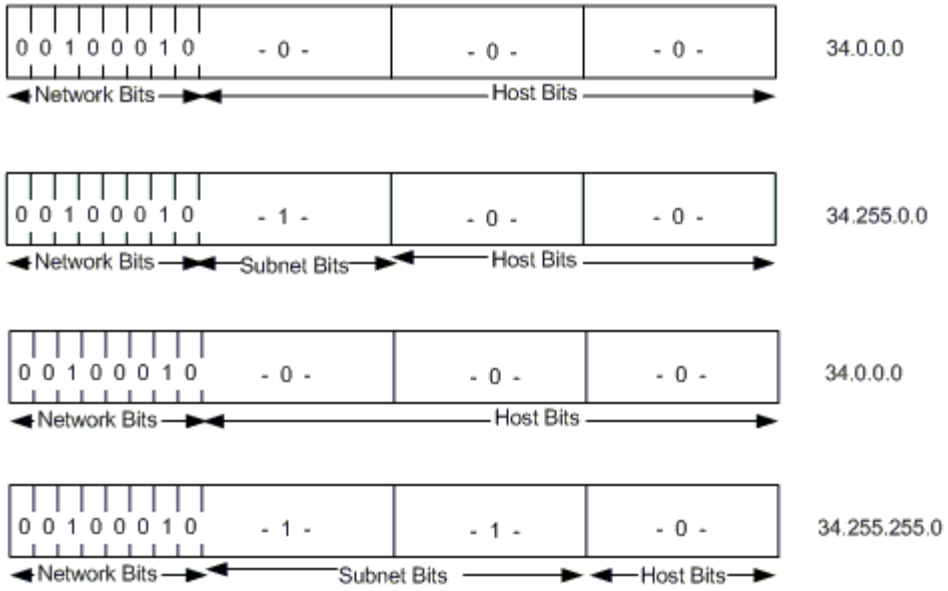

Traditionally, all subnets of the same network number used the same subnet mask. In other words, a network manager would choose an eight-bit mask for all subnets in the network. This strategy is easy to manage for both network administrators and routing protocols. However, this practice wastes address space in some networks. Some subnets have many hosts and some have only a few, but each consumes an entire subnet number. Serial lines are the most extreme example, because each has only two hosts that can be connected via a serial line subnet.

As IP subnets have grown, administrators have looked for ways to use their address space more efficiently. One of the techniques that has resulted is called Variable Length Subnet Masks (VLSM). With VLSM, a network administrator can use a long mask on networks with few hosts and a short mask on subnets with many hosts. However, this technique is more complex than making them all one size, and addresses must be assigned carefully.

Of course in order to use VLSM, a network administrator must use a routing protocol that supports it. Cisco routers support VLSM with Open Shortest Path First (OSPF), Integrated Intermediate System to Intermediate System (Integrated IS-IS), Enhanced Interior Gateway Routing Protocol (Enhanced IGRP), and static routing. Refer to IP Addressing and Subnetting for New Users for more information about IP addressing and subnetting. On some media, such as IEEE 802 LANs, IP addresses are dynamically discovered through the use of two other members of the Internet protocol suite: Address Resolution Protocol (ARP) and Reverse Address Resolution Protocol (RARP). ARP uses broadcast messages to determine the hardware (MAC layer) address corresponding to a particular network-layer address. ARP is sufficiently generic to allow use of IP with virtually any type of

underlying media access mechanism. RARP uses broadcast messages to determine the network-layer address associated with a particular hardware address. RARP is especially important to diskless nodes, for which network-layer addresses usually are unknown at boot time.

Routing in IP Environments

An "internet" is a group of interconnected networks. The Internet, on the other hand, is the collection of networks that permits communication between most research institutions, universities, and many other organizations around the world. Routers within the Internet are organized hierarchically. Some routers are used to move information through one particular group of networks under the same administrative authority and control. (Such an entity is called an autonomous system.) Routers used for information exchange within autonomous systems are called interior routers, and they use a variety of interior gateway protocols (IGPs) to accomplish this end. Routers that move information between autonomous systems are called exterior routers; they use the Exterior Gateway Protocol (EGP) or Border Gateway Protocol (BGP). Figure 6 shows the Internet architecture.

Figure 6 – Representation of the Internet Architecture

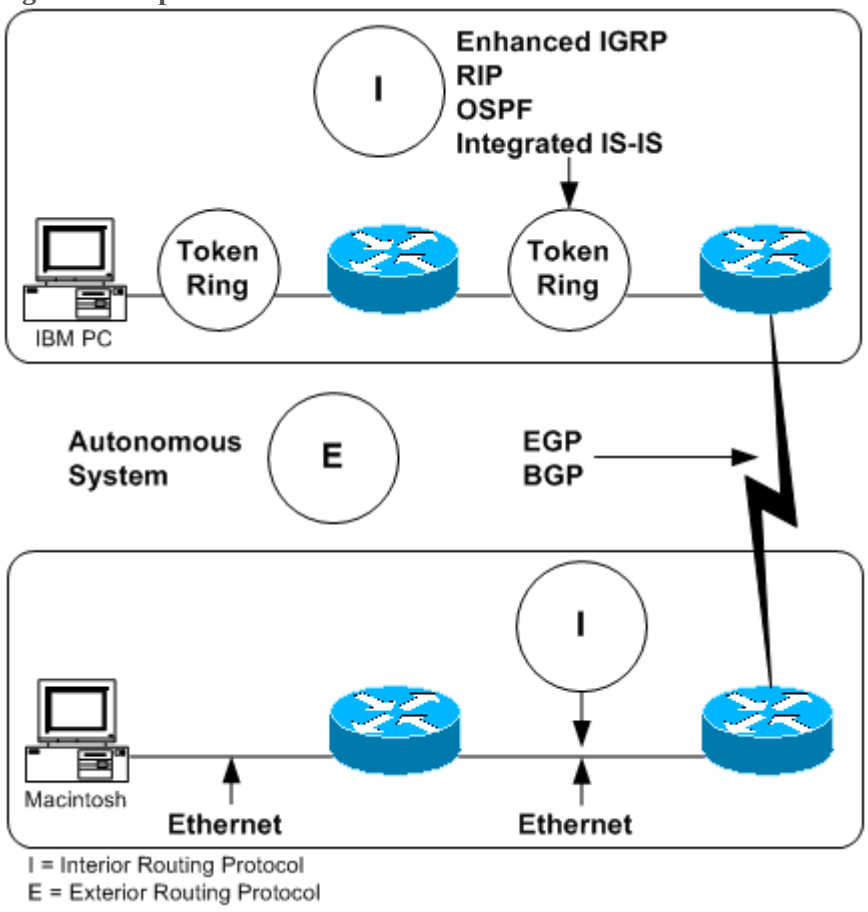

Routing protocols used with IP are dynamic in nature. Dynamic routing requires the software in the routing devices to calculate routes. Dynamic routing algorithms adapt to changes in the network and automatically select the best routes. In contrast with dynamic routing, static routing calls for routes to be established by the network administrator. Static routes do not change until the network administrator changes them.

IP routing tables consist of destination address/next hop pairs. This sample routing table from a Cisco router shows that the first entry is interpreted as meaning "to get to network 34.1.0.0 (subnet 1 on network 34), the next stop is the node at address 54.34.23.12":

```
R6-2500# show ip route
   Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
   D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
```

```
   N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
   E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
   i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
   ia - IS-IS inter area, * - candidate default, U - per-user static route
   o - ODR, P - periodic downloaded static route
Gateway of last resort is not set
          34.0.0.0/16 is subnetted, 1 subnets
O                 34.1.0.0 [110/65] via 54.34.23.12, 00:00:51, Serial0
   54.0.0.0/24 is subnetted, 1 subnets
C                  54.34.23.0 is directly connected, Serial0
R6-2500#
```

As we have seen, IP routing specifies that IP datagrams travel through an internetwork one router hop at a time. The entire route is not known at the outset of the journey. Instead, at each stop, the next router hop is determined by matching the destination address within the datagram with an entry in the current node's routing table. Each node's involvement in the routing process consists only of forwarding packets based on internal information. IP does not provide for error reporting back to the source when routing anomalies occur. This task is left to another Internet protocol—the Internet Control Message Protocol (ICMP).

ICMP performs a number of tasks within an IP internetwork. In addition to the principal reason for which it was created (reporting routing failures back to the source), ICMP provides a method for testing node reachability across an internet (the ICMP Echo and Reply messages), a method for increasing routing efficiency (the ICMP Redirect message), a method for informing sources that a datagram has exceeded its allocated time to exist within an internet (the ICMP Time Exceeded message), and other helpful messages. All in all, ICMP is an integral part of any IP implementation, particularly those that run in routers. See the Related Information