Notes

CTI 2500 Series PLC System

Application Note

IP Addressing and the 2572

The Origins of TCP/IP and the Internet

TCP/IP has Cold War origins, dating back to a networking experiment begun in the late 1960s under the auspices of the United States' Defense Advanced Research Projects Agency (DARPA). DARPA's "internetting project" called for developing technologies that would connect DARPA's computer network (ARPANET) with other networks located throughout the world, even if parts of the network would become damaged or possibly even destroyed as a result of war or civil insurrection. The most important invention that came from that research effort was the development of TCP/IP, which is the network protocol upon which Internet connectivity and the 2572 are based.

In 1986, the National Science Foundation (NSF) used the TCP/IP network protocol to create its own network (NSFNET), which allowed researchers to share data and access resources located on other remote networks. NSFNET served as the "backbone," or port of entry, for the U.S. portion of the Internet. The establishment of NSFNET was a watershed event for the Internet, because it provided the bandwidth and the network infrastructure needed to support the Internet's future growth. It is noteworthy that NSFNET initially provided backbone speed, or "bandwidth" of 56 kilobits per second, which is the same speed as many analog modems that are in use by consumers today.

During the early 1990s, the federal government began to privatize the backbone by transferring it over to a number of private companies, a process that was completed by April 1995. Since privatization has occurred, the backbone providers have been constantly upgrading network bandwidth capacity. Bandwidth upgrades are needed to support the increasing number of users coming on-line, and to accommodate adequately the bandwidth requirements of the increasingly larger proportion of network traffic devoted to the delivery of multimedia content associated with the World Wide Web.

IP Addressing

Basic Structure

Because TCP/IP networks are interconnected widely across the world, every machine on the Internet must have a unique address to make sure that transmitted data reaches the correct destination. Blocks of addresses are assigned to organizations by the Internet Assigned Numbers Authority (IANA). Individual users and small organizations may obtain their addresses either from the IANA or from an Internet service provider (ISP). The Internet Protocol uses a 32-bit address structure, shown as four 8-bit octets. The address is usually written in dot notation (also called dotted-decimal notation), in which each group of eight bits is written in decimal form, separated by decimal points. For example, the binary address:

11000011 00100010 00001100 00000111

is normally written in decimal as:

195.34.12.7

which is easier to remember and easier to enter into your computer.

Control Technology Inc.
5734 Middlebrook Pike, Knoxville, TN 37921-5962
Phone: 865/584-0440 Fax: 865/584-5720 www.controltechnology.com

Networks and Nodes

The 32 bits of the address are subdivided into two parts. The first part of the address identifies the *network*, and the second part identifies the *node* or station on the network. The dividing point between the network and node part of the address may vary depending on the address range and the application.

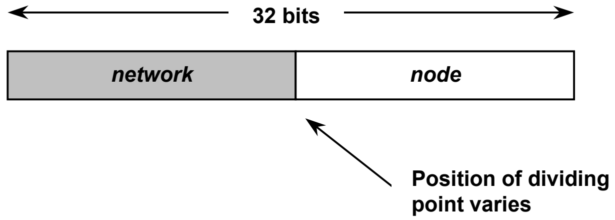

There are five standard classes of IP addresses. These address classes have different ways of determining the network and node sections of the address, allowing for different numbers of nodes on the network. Each address type begins with a unique bit pattern, which is used by the TCP/IP software to identify the address class.

Once the address class has been determined, the software can correctly identify the node section of the address. The three main address classes are illustrated below, showing the network and node sections of the address for each address type.

CLASS A Address

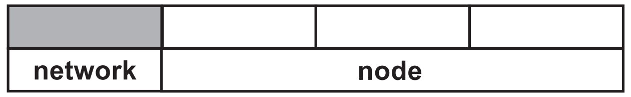

Class A addresses can have up to 16,777,214 nodes on a single network. They use an 8-bit network number and a 24-bit node number. Class A addresses are in the range:

1 . x . x . x to 126 . x . x . x

CLASS B Address

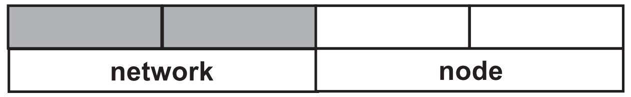

Class B addresses can have up to 65,354 nodes on a network. They use a 16-bit network number and a 16-bit node number. Class B addresses are in the range:

128 . 1 . x . x to 191 . 254 . x . x

CLASS C Address

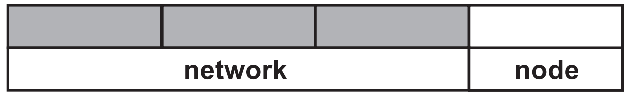

Class C addresses can have up to 254 nodes on a network. They use a 24-bit network number and an 8-bit node number. Class C addresses are in the range:

192 . 0 . 1 . x to 233 . 255 . 254 . x

in V-memory and execution of these command blocks from PLC logic. This method is beyond the scope of this Application Note. Refer to the 2572 Installation and Operation Guide Section 2.6 for a detailed discussion with examples.

2. Using network parameters stored in module EEPROM.

   Before you can automatically start the Network Server, the network parameters, including the module IP address, must be stored in EEPROM on the 2572. You can accomplish this task using a PC and a utility program from CTI.

   a. Attach the PC to Port 1 (RS-232) using a serial cable wired for RS-232. The cable that you use with TISOFT should work properly. *NOTE: Make sure that the communications parameters set for the 2572 match those of the PC and that the CAMP/NITP protocol is selected. (See section 2.4 of the Installation and Operation Guide.)*
   b. Place the diskette labeled *CTI 2572 Utilities* in a 3.5” diskette drive.
   c. Run the **IPSET** program from the diskette.
   d. Follow the instructions on the screen for establishing the network parameters. *NOTE: Ensure that the module EEPROM Write protect switch is off. (See Figure 13 of the Installation and Operation Guide.)*

You can also use the IPSET program to read the network parameters contained in the EEPROM. Complete instructions for using the IPSET program can be found in the IPSET.TXT file located on the 2572 Utilities diskette.

Once you have completed setting the IP address, you should power down the module and ensure that the Network Startup Option Switch is set to AUTOSTART. You may also wish to set the EEPROM Write Protect switch to ON. The new IP address will take effect when power is reapplied to the module. *NOTE: If you do not set the Network Startup Option switch to the* AUTOSTART *position and there is no PLC logic to set the network parameters, the IP address will* not *be set. The module* ACTIVE LED *will continue to blink.*

Using *subnetting*, the Class B address can be effectively translated into multiple Class C addresses. For example, the IP address of **172.16.0.0** is assigned, but node addresses are limited to 255 maximum, allowing 8 extra bits to use as a subnet address. The IP address of **172.16.97.235** would be interpreted as IP network address 172.16, subnet number 97, and node number 235. In addition to extending the number of addresses available, subnet addressing provides other benefits. Subnet addressing allows a network manager to construct an address scheme for the network by using different subnets for other geographical locations in the network or for other departments in the organization.

Although this example uses the entire third octet for a subnet address, note that you are not restricted to octet boundaries in subnetting. To create more network numbers, you need only shift some bits from the node address to the network address. For instance, to partition a Class C network number 192.68.135.0 into two, you shift 1 bit from the node address to the network address. The new netmask (or subnet mask) is 255.255.255.128. The first subnet has network number **192.68.135.0** with nodes 192.68.135.1 to 129.68.135.126, and the second subnet has network number 192.68.135.128 with nodes 192.68.135.129 to 192.68.135.254.

Table 1 lists the additional subnet mask bits in dotted-decimal notation. To use the table, write down the original class netmask and replace the 0 value octets with the dotted-decimal value of the additional subnet bits. For instance, to partition your Class C network 204.247.203.0 with subnet mask 255.255.255.0 into 16 subnets (4 bits), the new subnet mask becomes 255.255.255.240.

Note: The number 192.68.135.127 is not assigned because it is the broadcast address of the first subnet. And 192.68.135.128 is not assigned because it is the network address of the second subnet.

Table 1. Netmask notation translation for one octet

Number of bits	*Dotted-decimal value*
1	128
2	192
3	224
4	240
5	248
6	252
7	254
8	255

Table 2 displays several common netmask values in both the dotted-decimal and the masklength formats. CTI strongly advises that all nodes on a LAN segment use the same netmask for the following reasons:

- So that nodes recognize local IP broadcast packets. When a device broadcasts to its segment neighbors, it uses a destination address of the local network address with all ones for the node address. In order for this scheme to work, all devices on the segment must agree on which bits comprise the node address.
- So that a local router or bridge will know which addresses are local and which are remote.

Table 2. Netmask formats

Dotted decimal	*Masklength*
255.0.0.0	/8
255.255.0.0	/16
255.255.255.0	/24
255.255.255.128	/25
255.255.255.192	/26

Table 2. Netmask formats (cont.)

Dotted decimal	*Masklength*
255.255.255.224	/27
255.255.255.240	/28
255.255.255.248	/29
255.255.255.252	/30
255.555.255.254	/31

Private IP Addresses

If your networks are isolated from the Internet (for example, only between your two branch offices), you can assign any IP addresses to the nodes without problems. However, the IANA has reserved the following three blocks of IP addresses specifically for private networks:

10.0.0.0 - 10.255.255.255
172.16.0.0 - 172.31.255.255
192.168.0.0 - 192.168.255.255

CTI recommends that you choose your private network number from this list. Regardless of your particular situation, do not create an arbitrary IP address; always follow the guidelines explained here. For more information about address assignment, refer to RFC 1597, *Address Allocation for Private Internets*, and RFC 1466, *Guidelines for Management of IP Address Space.*

Basic Router Concepts

In general, the cost of providing network bandwidth is proportional to the data speed and the distance over which the network extends. Large amounts of bandwidth are provided easily and relatively inexpensively in a local area network (office, department and similar situations). However, providing the same high data speeds between two local networks that are physically distant may be prohibitively expensive. Because of this expense, high-speed local area networks (LANs) are usually interconnected by slower-speed links to form a wide area network (WAN). In order to make the best use of the slower WAN links, there must be a mechanism in place at each location for selecting data meant only for another location and sending it by the best available link. The function of selecting and forwarding this data is performed by a router.

What is a Router?

A router is a device that forwards traffic between networks based on network layer information in the data and on routing tables maintained by the router. In these routing tables, a router builds up a logical picture of the overall network by gathering and exchanging information with other routers in the network. Using this information, it chooses the best path for forwarding network traffic. Routers vary in performance and scale, number of routing protocols supported and types of physical WAN connections supported.

Routing Information Protocol

One of the protocols used by a router to build and maintain a picture of the network is the Routing Information Protocol (RIP). RIP is a distance vector protocol, meaning that all the decisions about which path to use are based upon a logical distance between source and destination networks. This distance is measured in “hops,” meaning the number of relaying routers in the path between the source LAN router and the router of the destination LAN. For example, the LAN of router A is considered to be 1 hop away. If router A can reach the network of router B by a direct WAN link to the network of router B, the network of router B is 2 hops away. If another network must be reached by calling router B and having router B forward the data, that network is n hops away, where n is the number of routers traversed by the data to get to the network farthest away. When there are multiple paths to a network, the path with the fewest number of hops is chosen and is regarded as the best path, and all other information about how to get to that network is discarded.

Using RIP, routers update one another periodically and check to see if there are any changes to be added to the routing table. An important consideration is the convergence time, or how long it takes for a change to the routing topology (such as a new node or a node failure), to be propagated throughout the entire RIP environment. To prevent this convergence process from being excessively long, RIP is limited to 15 hops maximum.

Address Resolution Protocol

An IP address alone cannot be used to deliver data from one device to another on a LAN. In order for data to be sent from one device on the LAN to another, you must convert the IP address of the destination device to its media access control (MAC) address. Each device on an Ethernet network has a unique Ethernet MAC address, which is a 48-bit number assigned to each device by the manufacturer. The technique that associates the IP address with a MAC address is known as address resolution, and IP uses the Address Resolution Protocol (ARP) to do this. If a device needs to send data to another station on the network and it does not already have the destination MAC address recorded, ARP is used. An ARP request is broadcast onto the network, and all stations receive and read the request. The destination IP address for the chosen station is included as part of the message so that only the station with this IP address responds to the ARP request and all other nodes discard it. The node with the right IP address responds with its own MAC address directly to the sender, providing the transmitting station with the destination MAC address needed for it to send the data. The IP address data and MAC address data for each node are held in an ARP table, so that the next time data needs to be sent, the address can be obtained from the address information in the table.

Domain Name Server

Many of the resources on the Internet can be addressed by simple descriptive names such as *www.controltechnology.com*. This addressing is very helpful at the application level, but the descriptive name must be translated to an IP address in order for a user to actually contact the resource. Just as each workstation maintains an ARP table to map IP addresses to MAC addresses, a domain name server (DNS) maps descriptive names of network resources to IP addresses. When a workstation needs to access a resource by its descriptive name, it first contacts a DNS to obtain the IP address of the resource. It can then send the desired message using the IP address. Many large organizations such as ISPs maintain their own DNSs and allow their customers to use them for address lookup. Presently, the 2572 does not use a domain name server. All client connections are made by the 2572 using specific IP addresses.

IP Configuration by DHCP

When an IP-based local area network is installed, each workstation must be configured with an IP address. If the workstations need to access the Internet, they should also be configured with a gateway address and one or more DNS server addresses. As an alternative to manual configuration, there is a method by which each device on the network can obtain this configuration information automatically. A device on the network may act as a Dynamic Host Configuration Protocol (DHCP) server. The DHCP server stores a list or pool of IP addresses, along with other information (such as gateway and DNS addresses) that it may assign to the other devices on the network. The 2572 does not presently support use of a DHCP server. Its IP address must be explicitly assigned, either by the PLC logic or by using the IPSET program and a serial cable to the module.

Setting the IP Address on the 2572

Two methods can be used to set the IP address and other network parameters on the 2572:

1. Using network parameters from PLC logic - this method has the benefit of tying the IP address to the PLC and not the module. In the event of a module failure and swapout, the process will continue to run without need for setup of the new module. It requires setup of command blocks

Two other classes of addresses exist, Classes D and E, which are used for multicasts (messages sent to many nodes) and for experimental use.

This addressing structure allows IP to uniquely identify each physical network and each node on each physical network.

For each unique value of the network portion of the address, the base address of the range (node address of all zeros) is known as the network address and is not usually assigned to a node. Also, the top address of the range (node address of all ones) is not assigned but is used as the broadcast address for sending a packet simultaneously to all nodes with the same network address.

Netmasks

In each of the above address classes, the size of the two parts (network address and node address) is implied by the class. This partitioning scheme can also be expressed by a *netmask* associated with the IP address. A netmask is a 32-bit quantity that, when logically ANDed with an IP address, yields the network address. For instance, the netmasks for Class A, B and C addresses are 255.0.0.0, 255.255.0.0 and 255.255.255.0, respectively. For example, the address

192.168.170.237

is a Class C IP address whose network portion is the upper 24 bits. When ANDed with the Class C netmask, as shown below, only the network portion of the address remains:

11000000 10101000 10101010 11101101 (192.168.170.237)

ANDed with

11111111 11111111 11111111 00000000 (255.255.255.0)

Equals

11000000 10101000 10101010 00000000 (192.168.170.0)

As a shorter alternative to dotted-decimal notation, the netmask may also be expressed in terms of the number of ones from the left. This number is appended to the IP address, following a slash (/), as “/n.” In the example, the address could be written as 192.168.170.237/24, indicating that the netmask is 24 ones followed by 8 zeros.

Subnet Addressing

By looking at the addressing structures, you can see that even with a Class C address, there are a relatively large number of nodes per network. Such a structure is an inefficient use of addresses because few networks have the maximum number of nodes permitted. Without some technique of “preserving” the unused addresses, we would soon run out of IP addresses. This problem is resolved by using a technique known as *subnet addressing*. Subnet addressing allows us to split one IP network address into smaller multiple physical networks known as subnetworks. Some of the node numbers are used as a subnet number instead. For example, a Class B address gives us 16 bits of node numbers translating to 64,000 nodes. Most organizations do not use 64,000 nodes on a single network, so there are free bits that can be reassigned. Subnet addressing makes use of those bits that are free, as illustrated below:

Subnetted CLASS B Address

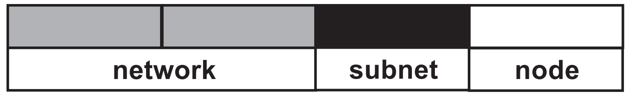
network
subnet
node