

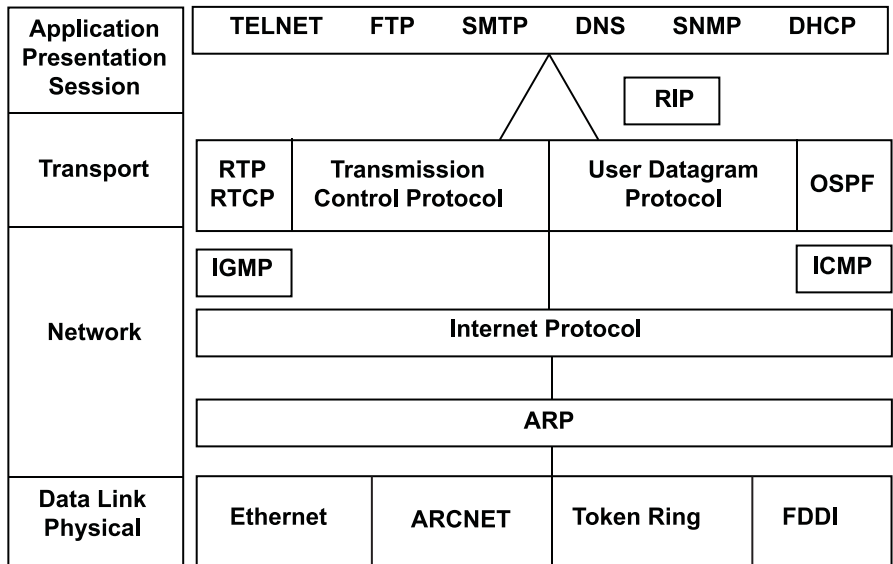
The ABCs of IP (Internet Protocol)

INTRODUCTION

The push to incorporate Industrial Ethernet or even “plain vanilla” Ethernet into control networks implies that by making that choice completes the selection process. As mentioned in a previous article, Ethernet II and IEEE 802.3 are strictly data link layer technologies which do not guarantee the delivery of messages over a network or between networks. Protocol stacks such as TCP/IP or SPX/IPX provide that functionality and without them Ethernet would be useless. With the immense interest in the Internet and the potential of attaching control networks to the Internet, the protocol stack of choice is TCP/IP because it provides the foundation for the Internet. This article addresses issues related to the IP portion of the TCP/IP stack as it applies to control networks.

THE TCP/IP STACK

Actually TCP/IP is a set of protocols defined by a series of RFCs (request for comments) that have evolved over the years. In general the Internet Protocol (IP) is used to route messages between networks and, therefore, properly resides at the network layer of the OSI Reference Model. Transmission Control Protocol (TCP) sits on top of IP and is used to guarantee the delivery of messages. Above TCP is the application layer. The services of the presentation and session layers of the OSI Reference Model are incorporated into the application layer. Therefore, the reference model for TCP/IP-based systems actually consists of only five layers.



Technologies such as Ethernet II and IEEE 802.3 reside at the lower data link and physical layers of the same model.

DATA ENCAPSULATION

The data sent over wires is represented as frames. An Ethernet II frame consists of a preamble, source and destination addresses, type field, data field and a frame sequence check field. You lump these fields into Ethernet header, data and trailer fields. The IP data sits above the data link layer and its data, called a datagram, is inserted into the data field of the Ethernet frame. The datagram has its own header and data fields but no trailer field. Above the IP layer is the transport layer where TCP and User Datagram Protocol (UDP) reside. Data from this layer is likewise applied to the data portion of the IP datagram. TCP applies segments while UDP applies datagrams. Both TCP and UDP have headers as well.

Figure 1. The TCP/IP stack is actually a set of protocols. IP resides at the network layer of the OSI Reference Model shown on the left.

Finally above the transport layer is the application layer which needs to insert its own data into the data portion of the transport layer as well as its own header.

This application data is simply referred to as data since there is no defined structure in terms of the TCP/IP stack. That is why if two application data structures are different, communication between these applications will not be effective even with strict adherence to TCP/IP standards.

The wrapping of data within the data field of the next immediate lower layer of the protocol stack is called encapsulation while the unwrapping of the same data at the receiving side is called demultiplexing.

Contemporary Control Systems, Inc. • 2431 Curtiss Street • Downers Grove, Illinois 60515 • USA

Telephone 1-630-963-7070 **Fax** 1-630-963-0109 **E-mail** info@ccontrols.com **Web** www.ccontrols.com, www.CTRLLink.com

Contemporary Controls Ltd • Sovereign Court Two • University of Warwick Science Park
• Sir William Lyons Rd. • Coventry CV4 7EZ UK

Telephone +44 (0)24 7641 3786 **Fax** +44 (0)24 7641 3923 **E-mail** info@ccontrols.co.uk

Web www.ccontrols.co.uk, www.CTRLLink.co.uk

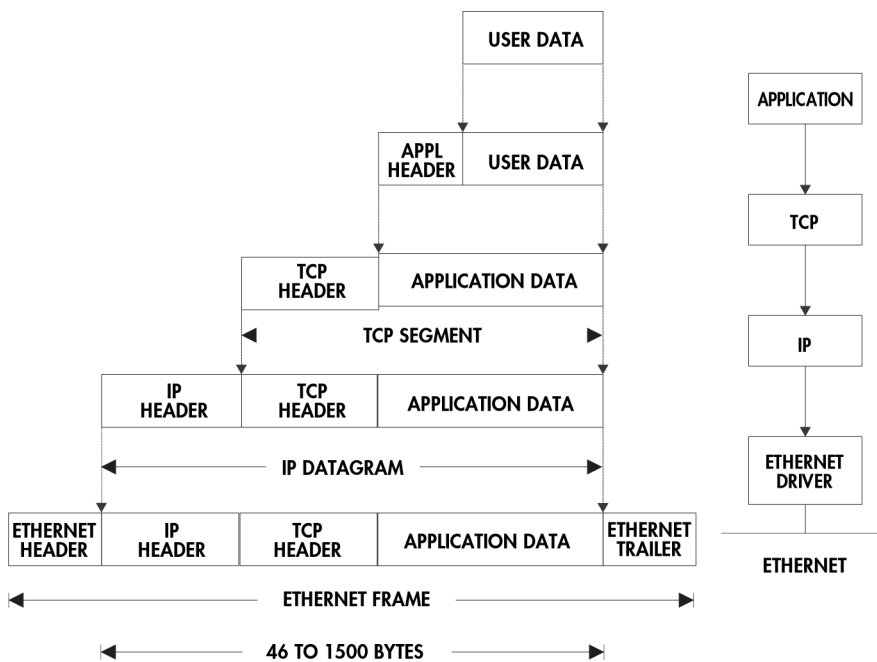


Figure 2. The wrapping of data into the data field of the next immediate lower layer is called encapsulation.

In order to reduce confusion on what is the actual data we will say that frames are sent over the data link layer. The IP sends out datagrams to the data link layer in the form of packets. A packet can be a datagram or a fragment of a datagram. The TCP sends segments while UDP sends datagrams. Finally, the application sends data. To further add to the confusion, the terms packet and frame are sometimes used interchangeably.

THE INTERNET PROTOCOL

The Internet Protocol provides the basic unit of data transfer, provides addressing, routing and fragmentation. The Internet Protocol resides at the network layer and sends and receives blocks of data called datagrams received from upper layer software. IP feeds these datagrams to its attached data link layer which sends and receives these datagrams as a series of packets. A datagram is analogous to a first-class letter sent in the Post. In general, it will reach its destination but there is no formal acknowledgement that the letter was received like there would be with either a registered or certified mail. IP utilizes a "best effort" or "connectionless" delivery service between source and destination

addresses. It is connectionless because there was no formal session established between the source and destination before the data was sent. Packets can be lost as they traverse the network or networks; thereby, corrupting datagrams. It is not the responsibility of IP to guarantee the delivery of messages and, therefore, IP is frequently termed an unreliable delivery service. That may be a little harsh of a criticism of IP, but it is the responsibility of the transport layer and not the network layer to guarantee end-to-end message delivery. IP is simply responsible for the addressing and routing of datagrams.

ROUTERS AND HOSTS

Unlike repeaters that operate at the physical layer and bridges that operate at the data link layer, routers operate at the network layer. A router is used to interconnect two networks together to form an Internet. An Internet is a general term used to denote a collection of networks. It is not to be confused with the Internet which is the public network that requires strict addressing standards in order for different systems to communicate. With a control network, we may want to keep it completely private and not connect it to the Internet or the corporate Internet (sometimes called an Intranet), but if

we do we will need a router. This is being mentioned here because IP is a routable protocol and routers are used to implement the protocol.

The end-to-end devices on the Internet are called hosts. If two hosts are on the same local network, then messages are routed directly involving no routers. If the two hosts are on different networks, a router must pass the message. This is called indirect routing.

IP ADDRESSING

The IP is responsible for source and destination addresses and its structure is defined in RFC 761. IPv4 is the most common version of addressing and it uses 32-bit addressing. The newer IPv6 calls for 128-bit addressing and was developed because the explosive growth of the Internet will soon deplete the inventory of possible 32-bit addresses. IPv6 will not be discussed here since there is ample confusion in simply discussing 32-bit IP addressing.

An IP address must not only address a particular host but a particular network as well. The IP address must not be confused with the Ethernet II address which is a 48-bit address sometimes called the MAC address. The MAC address is used to facilitate communication only at the data link layer. The IP address facilitates communication over networks and must be universally recognized even if the host is an Ethernet II node attached to a local area network or a serial port attached to a modem.

The format of the address is <netid, hostid> but is shown as one 32-bit address split up as four bytes. However, each byte is shown as a decimal number from 0 to 255. Therefore, an IP address is usually represented as XXX.XXX.XXX.XXX. This address can be shown as a binary or hexadecimal number as well, but the decimal-dot-decimal notation is the most popular. Therefore, the range of addresses is from 0.0.0.0. to 255.255.255.255.

An example of an address would be 128.8.120.5, but looking at the address it is hard to tell which is the network address and which is the host address.

There are five classes of IP addresses: A, B, C, D, E. Class D is for multicasting, a message from one host to many hosts, and class E is reserved for experiments. That leaves classes A, B and C which are the most important. These three classes break up the 32-bit address field into defined address ranges for the netid and hostid. You need to examine the very first byte of the IP address to determine the class. If the first bit of this byte is a '0' then this is a class A address. In a class A address the first byte identifies the network and the remaining three bytes identifies the host. That means you can have 16, 277, 214 hosts for every network!

identifies the host. This provides a reasonable 254 hosts.

Class D and class E addresses can be identified in the same way. A class D address has a leading bit pattern of "1110" while a class E address has a leading bit pattern of "11110."

There are also other reserved addresses. Regardless of class, a host address of all 1s is reserved for a broadcast message to all hosts on that network while a host address of all 0s is reserved to mean "this network." Network address 127 is also reserved and is used for loopback testing. This effectively wastes 16 million possible host addresses. Network address 0 is reserved as well.

If the control network is to become part of the public Internet then strict adherence to the class addressing rules must be followed. Usually these addresses will be

Class A :	1-127
Class B :	128-191
Class C :	192-223
Class D :	224-239
Class E :	240-254

Figure 4. The class of an IP address can be quickly identified by observing only the first byte.

IP HEADER

IP transmits and receives datagrams. Within the datagram is a header and the data portion of the datagram. The minimum size of the IP header is 20 bytes consisting of five 32-bit words. The first three words provide control information while the remaining two words provide address information. An optional field can follow the address information. The information in the header is as follows:

Version: A four-bit field identifies the IP version. A 4 identifies IPv4 while a 6 identifies IPv6.

Header Length: A four-bit field indicates how many four-byte words are in the header. The header length cannot exceed 60 bytes; thereby, allowing 40 bytes for options.

Type of Service: Of the eight-bit field only six bits are used. The Delay bit indicates the datagram should be processed with low delay by the router. The Throughput bit requests high throughput while the Reliability bit requests high reliability. There are three other bits to indicate precedence. These bits are set at higher layers of the protocol stack and are suggestions given the router. This looks like a nice feature for control networks since control networks require low delay and high reliability. However, it is not clear that routers even look at these bits. It appears that this was a feature with great promise but never really implemented. This is to be rectified in IPv6.

Total Length: The total length of the datagram including the header cannot exceed 65,535 bytes. This 16-bit field is for the datagram itself and not the packet length in the data link layer.

Address Identifier	Network Address	Host Address
Class A		
0	7 bits of network address	24 bits of host address
	First byte	Last three bytes
Class B		
10	14 bits of network address	16 bits of host address
	First two bytes	Last two bytes
Class C		
110	21 bits of network address	8 bits of host address
	First three bytes	Last byte
Class D		
1110	Multicast address in the range of 224.0.0.0 - 239.255.255.255	
Class E		
11110	Class E - Reserved for future use	

Figure 3. Address classes define the split between network and host IDs.

If the first two bits of the first byte are a "10," then this is a class B address. With class B addresses the first two bytes identify the network and the remaining two bytes identify the host. This provides a slightly more reasonable 65, 534 host addresses.

If the first three bits of the first byte are a "110," then this is a class C address. With class C addresses the first three bytes identify the network and the remaining byte

issued by the corporate network administrator or by an Internet Service Provider (ISP). But what if the control network is to become strictly a private network? Cannot any addressing scheme work? Yes, any address scheme could work but there is even an RFC guideline for this situation. According to RFC 1918, only non-routable IP addresses should be used. These addresses, which a router will not pass, are as follows:

10.0.0.0 to 10.255.255.255

172.16.0.0 to 172.31.255.255

192.168.0.0 to 192.168.255.255

If this datagram is larger than the maximum packet length that can be sent, the datagram will need to be fragmented into manageable successive packets. In this case the total length field will represent the length of the fragment sent and not the length of the original datagram.

Ideally you would like to restrict datagram size to packet size in order to avoid fragmentation. With Ethernet II the maximum packet size is 1500 bytes. This is called its Maximum Transmission Unit (MTU) and within a private or local network the MTU is known and can be adhered to. The problem occurs between networks. Intermediate networks may have a lesser MTU requiring

If fragments are to be sent, it is necessary to load in the fragment offset. Notice that with every fragment, the IP header is resent with just a slight modification. The fragment offset will change on every fragment and possibly along with one flag bit. Fragments must be sent in eight-byte multiples because there are only 13 bits available for identifying fragments and datagrams can be 64 KB in length. For example, if the first fragment is 1024 bytes long, the fragment offset of the next fragment will indicate that the accompanying fragment begins the 1025th byte of the original datagram. With knowledge of the datagram identifier, fragment offset, the source IP address and the fragments themselves, the complete datagram can be reassembled by the receiving host even if the fragments are received out of order. That is the true strength of the IP. Packets can take different routes to the intended destination and still be reassembled into the original datagram.

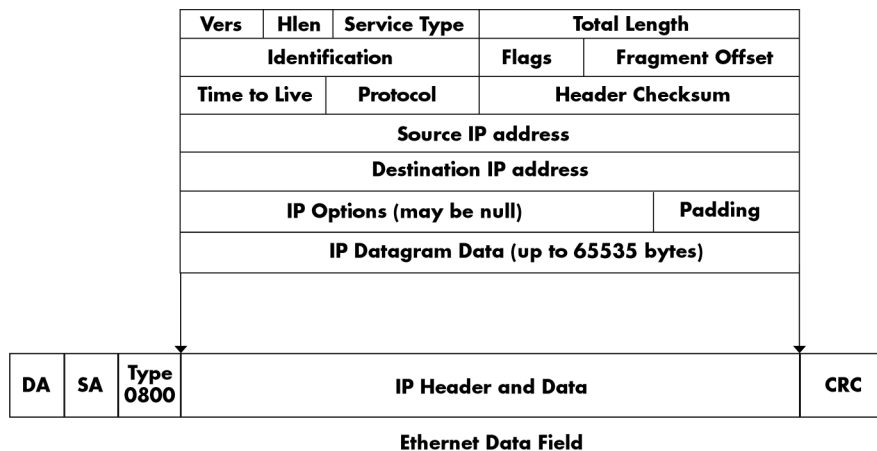


Figure 5. The IP datagram consisting of a header and data is inserted into the Ethernet data field.

Datagram Identification: A unique 16-bit identifier assigned by the host will accompany the datagram. This is necessary in order for the receiving host to reassemble fragmented datagrams. All fragments will contain the same datagram identifier.

Flags: Three bits are reserved for flags but only two are used. The Don't Fragment bit tells the router not to fragment the datagram. If this cannot be done, an error message is returned. The More Fragments bit is used in the fragmentation process. A 1 means that the datagram being sent is actually a fragment of a larger datagram. A 0 means that either the datagram is not fragmented (first and only datagram) or it's the last fragment. Receiving hosts need this information in order to reassemble fragments.

Fragment Offset: Thirteen bits are used to indicate which fragment is being sent. Fragmentation is the process of breaking up large datagrams into manageable packets.

the router to fragment the original message even though it was originally sent unfragmented. The router does the fragmentation on his own (as long as the datagram was not marked as "do not fragment") and the fragments must be recombined at the destination host. Routers do not recombine fragments.

The default MTU is 576 bytes and all routers must be able to handle that size transmission. By restricting the datagram to 576 bytes, it will never need to be fragmented. Of course, that puts an undue restriction on the Ethernet II network since packets can be as long as 1500 bytes. So for local networks set the maximum datagram size to the local network's MTU. If the datagram is to be sent beyond the local network, set the maximum datagram size to 576 bytes. For control networks, fragmentation may never be an issue since control information packets are usually short, not exceeding 256 or 512 bytes. Fragmentation should be avoided since it increases data latency and increases the chances of a corrupted datagram since multiple packets must be sent per datagram.

Time to Live: This eight-byte field is strictly used by the routers to prevent a datagram from a faulty transmission sequence to endlessly circulate around an Internet. Originally, the unit of measure was seconds because it was believed that it would take a router one or more seconds to process a datagram from its queue. Once the datagram was processed, the router would decrement this field by the amount of time that occurred. However, in practice, modern routers are much faster than early routers and usually process the datagram within a second, but only decrement the field by one (the minimum amount). Therefore, the field has come to be treated as a hop counter. A hop being an instance of a datagram being processed by a router. The originating host sets the Time to Live field and each router decreases it by one. If a router decrements the count to zero, it will discard the datagram and inform the originating host that the datagram failed to reach its destination.

Protocol: The eight-bit protocol field informs the upper layer protocol that the received datagram is for its use. Usually, the upper layer protocol is TCP or UDP, but there are other protocols as well that

could be sending and receiving data. The protocol field provides this distinction.

Header Checksum: The complete IP header is checked for integrity with the aid of the 16-bit header checksum. The originating host applies the checksum and all routers check the header for integrity and regenerate a new checksum when the datagram is resent. A new checksum is required since the Time to Live field would have been changed by the router. Finally, the checksum is again reconfirmed by the receiving host.

Source/Destination Address: The 32-bit source and destination addresses are included in the header. These are the IP and not MAC addresses.

IP Options: There may be no options in which case this field is null or there can be options usually intended for router use only. The option fields must be at least 32-bits in length and must be padded to that amount if shorter.

What Defines TCP/IP?

The TCP/IP stack and its associated protocols are described in Request for Comments (RFCs). There are about 2700 RFCs in existence and unlike many industrial control standards these are free! You can simply download them from the Internet. One possible location is <http://www.ietf.org>. Which RFCs do you need? Matthew Naugle, in his book called *Illustrated TCP/IP*, suggests as mandatory reading RFCs 1122, 1123, and 1812. These will give a good overview but you can always seek out others one you find an index. You can also author your own RFC by following the instructions and format in RFC 1543. Most of the RFCs originate from the working groups (WGs) of the Internet Engineering Task Force (IETF). Some RFCs obsolete prior RFCs. If your control strategy is based upon the TCP/IP protocol, it is recommended that your document which RFCs are important to your system.

ARP

As mentioned before, the IP routes datagrams between source and destination addresses in the form of packets over a data link layer. The data link does not understand datagrams nor does it understand IP addresses. It does know, however, its own MAC address and knows how to communicate to other MAC addresses when told to do so. Somehow we need to inform each host what IP address its MAC address or physical address has been assigned, and we need to inform the same host all the other physical address assignments on the local network in order to have communication.

Usually, the host IP address-physical address assignment is stored in non-volatile memory or in a file. Using a 32-bit DIP switch for assignment is not practical. Sometimes a serial port on the device is used for programming the IP address, but once programmed, all other hosts on the local network must still need to learn the assignment.

The Address Resolution Protocol (ARP) is used for learning physical address assignments. ARP has its own structure and does not use that of IP. ARP directly communicates to the data link layer and, therefore, must be aware of the various types of network adapters that are available.

When a host needs to send a datagram to another host on a local network, it first checks its ARP table to determine the physical address for that IP destination address. If one is found, the datagram transmission proceeds. If none is found, an ARP request is made. An ARP request consists of a broadcast message to all hosts on the local network.

Within the ARP request is the originator's IP and physical addresses as well as the requested IP address.

REFERENCES

Illustrated TCP/IP, Matthew Naugle, 1998, Wiley Computer Publishing

Practical Networking with Ethernet, Charles E. Spurgeon, 1997, International Thomson Computer Press

International Standard ISO/IEC 8802-3 ANSI/IEEE std 802.3, 1996, The Institute of Electrical and Electronic Engineers, Inc.

TCP/IP Clearly Explained, Pete Loshin, 1997, Academic Press

TCP/IP Illustrated, Volume 1, The Protocols, W. Richard Stevens, 1994, Addison-Wesley Publishing Company

Since it is a broadcast message, all hosts have the opportunity to learn the IP address and physical address pairing of the requestor which can be appended to that host's ARP table. Only the host with the requested IP address responds to the ARP request by providing its IP address and physical address pairing. This message is sent as a unicast message back to the requestor. Once the physical address is known by the requestor, the datagram can be sent.

SUMMARY

The IP is responsible for the end-to-end delivery of datagrams over an Internet. It also provides host and network addressing and the means for fragmenting datagrams into manageable packets. IP is a routable protocol and much of its complexity is due to its ability to route packets directly within a local network or indirectly through routers. Routers are not ideal for a control network since they reduce determinism and increase data latency. Still to accept TCP as a transport layer for an Ethernet control network requires acceptance of IP as well. By understanding the limitations of IP, a control network can still be designed using the TCP/IP family of protocols. This is especially true if the control network is restricted to that of a private or local network.