VOLUME 3 • ISSUE 1

A TECHNICAL SUPPLEMENT TO CONTROL NETWORK

ESSENTIALS

Understanding Ethernet Switches and Routers

This extended article was based on a two-part article that was written by George Thomas of Contemporary Controls and appeared in the February and March 2011 issues of InTech Magazine — an ISA publication.

When you go to a computer store to purchase a device that will access the Internet they will try to sell you a *router*. In most instances, the router will come with a built-in *switch* (shortened term for *switching hub*) so that you can connect several Ethernet devices to just one device. So what is the difference between an Ethernet router and an Ethernet switch? The long answer to that question requires an examination of the Open Systems Interconnection Model (OSI) which is frequently used to explain how communication networks operate.

Ethernet and the OSI Model

In Figure 1 you will see the seven-layer model with each layer providing a unique service. Communication between two stations begins at the *Application* layer with the sending station initiating a message to a receiving station via a common *medium*. With respect to this model, Ethernet provides services at the *Physical* and *Data Link* layers through the use of bridges and repeaters. An Ethernet switch is classified as a bridge and therefore operates at the data link layer while routers operate at the *Network* layer. Let's try to understand why.

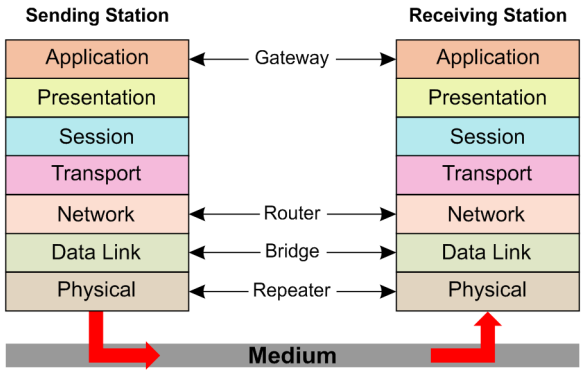

Figure 1 — The OSI Model

The lowest layer is the physical layer that defines the basic signalling on the medium. Ethernet transmits *symbols* representing logic "ones" or "zeros" across the medium to another station that decodes the symbols to extract the data. Although Ethernet will operate with coaxial cable as the medium, modern Ethernet networks incorporate twisted-pair cabling. If the path is too long, a repeater can be used to extend distance. If fibre optic cable is preferred, media converters can be used. If multiple devices need to share the connection, a repeating hub (commonly called just a hub) is used. All three of these devices reside at the physical layer because they do nothing more than process symbols on the medium.

One layer above the physical layer is the data link layer. Ethernet is a local area network technology with end stations assigned unique 48-bit addresses. These addresses are called *media access control* (MAC) addresses. When data is to be sent from one Ethernet station to another, the data is first arranged in *frames* as shown in Figure 2. The destination and source addresses are appended so that the intended station knows that it is to receive the message and who sent it. Other parts of the frame include the *Preamble* which alerts the receiving station that a frame is coming, a *Type or Length* field that identifies either the type of data or length of the data field, the *Data* field itself and the *Frame Check Sequence* used

Ethernet Frame					
64 bits	48 bits	48 bits	16 bits	368 to 12000 bits (46 to 1500 bytes)	32 bits
Preamble	Destination Address	Source Address	Type or Length	Data	Frame Check Sequence

Figure 2 — An Ethernet Frame

to verify the integrity of the frame. The payload of the frame is the actual data. Everything else is overhead.

Carrier Sense Multiple Access with Collision Detection (CSMA/CD)

As an Ethernet station prepares to send a frame, it first listens to the medium to verify that a clear channel exists. If silence is sensed, it transmits its message and then waits a determined amount of time, called the slot time. The slot time is used for detecting a collision due to another station transmitting at the same time. If no collision is sensed, the sender assumes a successful transmission. If a collision has occurred, the sender refrains from transmitting again for an amount of time based on a backoff algorithm that incorporates randomness. An early criticism of Ethernet was that this probabilistic approach to media access control was not conducive to real-time systems. There are other problems with CSMA/CD.

- All CSMA/CD stations must reside within one *collision domain* to ensure that all stations will detect a collision between stations located at the farthest points. This limits the geographic distance of the Ethernet network.
- Stations residing in the same collision domain can detect all transmissions but only receive those addressed to them. All stations can transmit but not at the same time. This is called half-duplex operation or *Shared Ethernet.*

Breaking Up the Collision Domains for Higher Performance

A switching hub was introduced to avoid the problems of Shared Ethernet. A switching hub is much different from a repeating hub. A port on a switching hub appears to an end station as another end station except that it does not consume a MAC address. To an attached end station, the switch port appears as the only other station within the collision domain. This is how it works.

Assume station A is on port 1 of an eight-port switch and station B is on port 2 as shown in Figure 3.

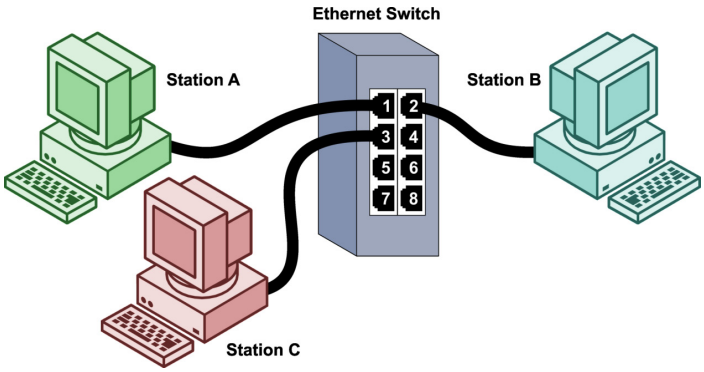

Figure 3 — Using an Ethernet Switch

Station A sends a message to station B. Switch port 1 reads the entire frame into its internal input buffer and forwards it to port 2's output buffer which then transmits the entire frame to station B. So what is the advantage?

- The switch has effectively created two collision domains — each appearing as a two-station link. With only two stations, collisions can be avoided altogether by creating a full-duplex link which potentially doubles throughput.
- With a full-duplex link there is no collision domain — thus, distance is limited only by cable losses. Fibre optic distances are no longer limited by the collision domain and can be much greater than Shared Ethernet lengths. Without a concern for a common collision domain, switches can be cascaded at will.
- With separate collision domains on each port, each port can operate at different data rates allowing for the mixing of data rates within the same switch.

Another advantage to using a switch is its ability for simultaneous messages within its switch fabric. When a transmission is received on a particular switch port,

the source MAC address of the sender is stored in the database of the switch. Using this *learning process*, shown in Figure 4, the switch determines on what port a station can be reached. Assume, in Figure 3, that station A sends a message to station C which has been attached to port 3, but the switch does not know how to reach station C. The switch will *flood* the same message to all ports. When station C eventually replies, the switch will learn that station C is on port 3 so that future flooding will not be necessary. Now station A sends a message to station B, but the switch already knows that station B can be reached on port 2 by using a *lookup* process so only port 2 will transmit the message. The other ports do not need to pass the message because it was only directed to station B. This frees up the other ports to pass unrelated messages without a concern for stations A and B's traffic. This greatly improves throughput over Shared Ethernet which requires that only one message can pass through a hub at any one time.

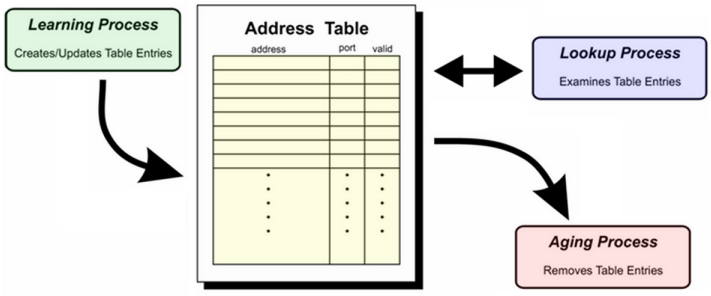

Figure 4 The Learning Process

Now assume the cable on port 1 is moved to port 4. If station A does not initiate a transmission, the switch will still believe station A can be reached on port 1. For this reason, all learned addresses must be *aged* by clearing out the database periodically. Eventually, the pairing of station A with port 1 will be cleared and transmissions intended for station A will be flooded to all ports, including port 4. Station A will now receive the flooded message and its response will allow the switch to learn its new location.

Modern switches have two more interesting features — Auto-negotiation and Auto-MDIX. With Auto-negotiation, the data rate and duplex for link partners is negotiated during initial connection. If the end station and the switch port can operate at either 10 Mbps or 100 Mbps at either half- or full-duplex, the negotiation process will select higher performing 100 Mbps full-duplex. With Auto-MDIX, either a straight-through or crossover cable can be used between an end station or switch port or between two switch ports.

Stepping Up to Routers

In the beginning of this article we used the example of visiting a computer store to purchase a device that will access the Internet and we noted that they will try to sell us a *router*. In most instances, the router will come with a built-in *switch* so that you can connect several Ethernet devices using just one device. So again, what is the difference between an Ethernet router and an Ethernet switch? We will refer back to the Open Systems Interconnection Model.

Revisiting the OSI Model

In Figure 1 you will see the seven-layer model with each layer providing a unique service. As we mentioned before, Ethernet provides services at the *Physical* and *Data Link* layers through the use of bridges and repeaters. The rules of Ethernet are restricted to a single local-area-network (LAN). If we have a collection of interconnected LANs, this is called an *inter-network.* Communicating between LANs within an inter-network requires *routers* which operate one layer above that of a switch at the *Network* layer. The most famous of inter-networks is the Internet with the rules for communication being defined by the *Internet Protocol* (IP).

It is not necessary that routers support the Internet Protocol, but this is the most common protocol used by routers so we will use this in our discussion. In Figure 5 you will see a collapsed seven-layer model which is called the Internet Model. The only difference is that the functions of *presentation* and *session* are lumped into the *transport* layer. The transport layer provides end-to-end communications between applications with the *Transmission Control Protocol* (TCP) being the one used in the Internet Model. The middle layer is the network layer which is involved with *host addressing* and *fragmentation* with the most common addressing scheme called IPv4.

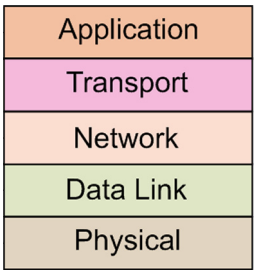

Figure 5 — The Internet Model

IPv4 Header and Datagram

Figure 6 shows an *IPv4 header and datagram* encapsulated into an Ethernet frame. The Ethernet frame is shown with a destination address (DA) and a source address (SA) as we would expect. The Type field has a hexadecimal 0800 indicating the Ethernet data field contains an IP *packet* or a portion of an IP packet. A Cyclic Redundancy Check (CRC) completes the frame. When we refer to the Internet Protocol we talk in terms of packets. With Ethernet communication we talk in terms of frames. Ethernet frames pass without issues through Ethernet switches, but it is the data inside Ethernet frames called packets that hosts and routers on the Internet respond to. Switches do not understand the meaning of packets — they just understand frames.

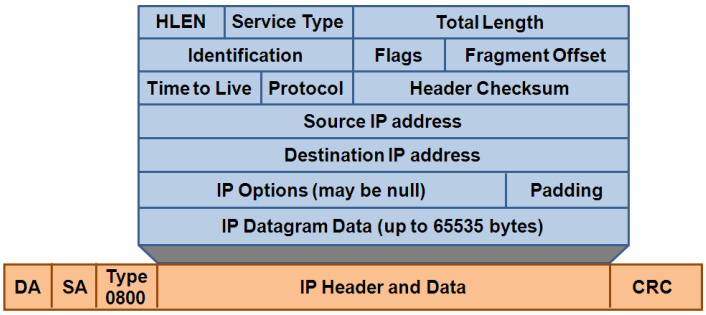

The version field will contain a 4 to indicate IPv4.

Figure 6 — IPV4 Header and Datagram

Figure 6 can be very confusing and it does not get any easier as we move up the Internet Model. We will just point out some fields within the packet that are the most interesting.

With the IPv4 (IP protocol version 4), a new addressing scheme is introduced to identify hosts on the Internet. There is a *source IP address* and a *destination IP address* just like there are source and destination Ethernet addresses so what is the difference? IP addresses are 32-bit long addresses while Ethernet addresses are 48-bits long. In the Internet world, a device does not have to be an Ethernet device in order to have an IP address. While an Ethernet device manufacturer supplies a unique 48-bit address to the product it sells, IP addresses are assigned using rules established by the *Internet Engineering Task Force* (IETF). IP addresses are designated as being *private* and *public*. Public IP addresses must be unique across the Internet, but private addresses can be duplicated because they will never reach the Internet — thanks to IP routers that block them from appearing on the Internet. Switches would not restrict IP addresses from appearing and that is one reason why routers are used to access the Internet and not switches.

Other fields that should be pointed out within the IP packet are the *IP Datagram Data*. Our main purpose is to send data between hosts and this data is referred to as a *datagram*. Unlike the transport layer of the Internet Model which is involved with end-to-end delivery of data, the network layer is only required to make its “best effort”. There are no acknowledgements that a packet sent by a host is received by another. If the packet is too big to fit into one Ethernet frame, it is split up into *fragments* requiring reassembly at the receiving end.

Ethernet Switch-Router Combination

Figure 7 shows a four-port Ethernet switch as part of an IP Router. Although there are a total of five Ethernet ports on this combination Ethernet switch-router, the five ports are not peers to one another. Certainly, the four clustered ports are peers — they are all part of the internal Ethernet switch and no one port on the switch has precedence over another. The switch-router is divided into two halves — the LAN-side and the WAN-side. The Ethernet switch resides on the LAN-side or what is sometimes called the private side. Remember that Ethernet is a local-area-network technology so Ethernet equipment deployment is restricted to either a work group, a building, process line or possibly a campus. Workstations, printers, servers and automation equipment can attach to any of these switch ports. If more LAN ports are needed, external switches can be cascaded to any switch port on the switch-router.

The one remaining Ethernet port is not part of the switch fabric. It is called a WAN port for *wide-area-network*. It does not have to be an Ethernet port but we will use an Ethernet port in our example. This single WAN port is considered to be located on the public side of the switch-router because this is the port that gains us access to the Internet. Between the LAN and WAN sides is logic that controls the routing of messages between the two sides. How can a single Ethernet port connect to the Internet? In our example we are attaching the Ethernet port to a cable modem although a Digital Subscriber Line (DSL) modem

could be used as well. At the far-end of the modem connection is an Internet Service Provider (ISP) that functions as a gate-keeper to the Internet.

Obtaining IP Addresses

In our example of Figure 7, the hosts on the private side have been given private IP addresses. These addresses can be statically set in the hosts or the hosts can receive them dynamically using a process called Dynamic Host Configuration Protocol (DHCP). In the dynamic case, the router functions as a DHCP server providing addresses in a preselected range while the attached hosts function as DHCP clients requesting IP addresses. However, the router needs a WAN side address as well and usually obtains a public IP address from the ISP using a similar process. Once addressing is established, connected hosts on the LAN-side can have access to the Internet through the switch-router. The translation of the private addresses to that of a public address is another function of a router.

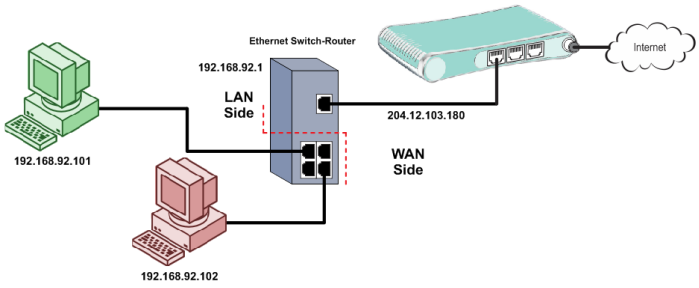

Figure 7 — Public and Private IP Addressing

Summary

Ethernet switches and routers have distinct purposes and when the two devices are included in the same enclosure — like a switch-router — it is sometimes difficult to separate the two functions. Think of Ethernet switches operating at layer 2 of the OSI model and routers operating one layer up at layer 3. At layer 2, we are concerned about the rules within one local-area-network while at layer 3, the rules for operating in an inter-network result in more complexity.