

Network Hubs vs. Switches

In the recent past, many industrial control networks relied heavily on proprietary networks to transmit data between PLCs. With the emergence of SCADA and HMI software, the desire to communicate between PLCs, HMI nodes, and other plant floor devices has begun a move towards a more open networking solution. Ethernet, with its easy accessibility in the commercial market and its open, multi-protocol ability, is showing promise as the communication network of choice. As people began to implement Ethernet, however, they realized that network topology and component selection play a major role in the performance and availability of the network. Although Ethernet is an extremely fast and inexpensive communication platform, it is not designed with packet traffic control systems that are inherent in proprietary control networks. Industrial Ethernet users often are concerned that if they connect plant-floor Ethernet nodes on the same shared Ethernet system with their office LAN, they run the chance of having an office event, such as a network back-up, create havoc and reduce the speed of their plant-floor control system. Understanding more about how Ethernet works can help industrial network planners make a more informed decision about the network devices and topologies to implement in their industrial Ethernet networks.

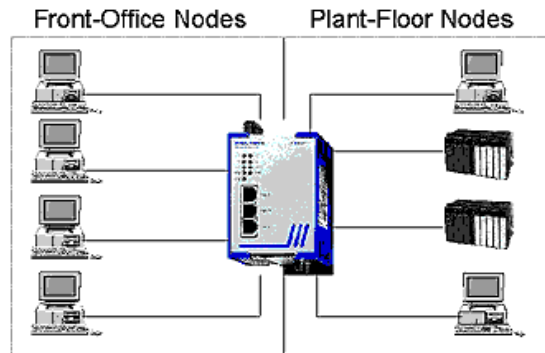
When an Ethernet network node wants to communicate, it first listens to its network connection to determine if there is any traffic currently on the network. If there is no traffic at the instant that the node checks the network, then the Ethernet node will transmit its Ethernet packet. Just like someone picking up a telephone and realizing that the line is in use by someone else in the house, the Ethernet node will delay communication until the network is free of traffic. The Ethernet node that is interested in communicating will check the Ethernet link until the line is free and then transmit when an open line is sensed. In many proprietary networks, traffic is controlled by the passing of a signal between the nodes on the computer that allows each node to transmit only during a controlled window of time. Ethernet has a "wild-west" multiple access architecture, justified in large part by the terrific speed and bandwidth available to the Ethernet network. However, in the multiple access model, two nodes which check the line and see no traffic can simultaneously transmit a packet, creating a collision of the two packets and a failed transmission. The collision of the two packets generates a signal that is recognized by all of the Ethernet devices on the shared network. The collision of the two packets causes all the nodes on the shared network to halt communication and wait a brief period of time before beginning to transmit again. Obviously, as the total amount of network traffic and/or nodes increases, so does the opportunity for collision. An important term to remember is the term "collision domain." A collision domain is simply a collection of devices on a shared Ethernet network, all of which are connected to the same shared, unregulated Ethernet network.

Let's consider a real-world example, illustrating in more common terms the multiple-access principle of Ethernet. Imagine a small business made up of five workers in an office. When the business opens, the owner decides to install a phone on each worker's desk, but only one shared phone line (Multiple Access Collision Domain.) The owner is pleased that he only has to pay for one line, and in the beginning all goes well. In the event that one of the workers in the office wants to dial out, he would pick the phone up to see if any of his co-workers were talking. If they were, the worker would put the receiver down, and try again until the line was free (Carrier Detect.) Occasionally, though, two workers would simultaneously pick up the phone, hear the dial tone, and begin to dial. The simultaneous dialing would cause neither person's phone number to be dialed, and both would put the phone down, wait a while, and try again (Collision Detect.) This is analogous to the Ethernet multiple-access principle. Any node in a collision domain has the ability to check the party line at any time and is not limited as to number of phone calls it can make or when to make a phone call. As additional nodes are placed on the network or the amount of data generated by nodes increases, the network can become increasingly burdened and slow overall communication rates.

Early plant-floor Ethernet networks often attempted to tie plant-floor equipment into the same Ethernet hubs that were handling front-office workstations. A hub is simply a multi-port broadcast device. It takes whatever comes in any port and broadcasts it out all the other ports. Even if two hubs are interconnected, you basically end up with one big collision domain, with all traffic shared. As network nodes are added or traffic increases, every node in the collision domain has a greater chance of slowing communication or having a collision. Additionally, since Ethernet nodes currently do not differentiate between the relative importance of Ethernet packets, it is possible for non-essential traffic on the network (perhaps people backing-up their computers to


the network server or printing a large document across the network) to slow or collide with essential traffic (such as inter-PLC communication, or HMI polling.)

Single Network



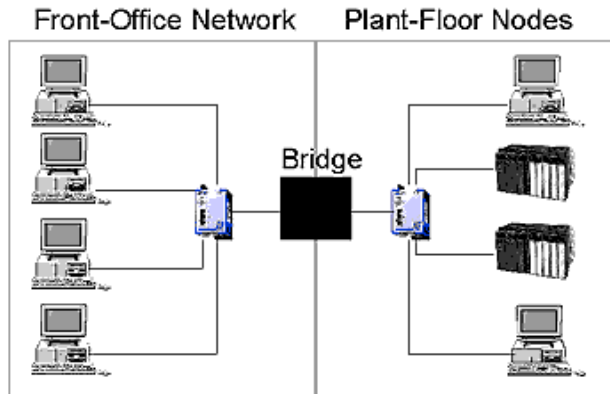
The first solution to this problem was simply to create a separate control network from the office LAN network. By having a separately-wired system, the plant floor could be assured of being immune from office LAN shutdowns or slowdowns. This philosophy becomes a problem when people want to share easily the information that is created in the plant network with office network nodes. If two networks are completely physically isolated, the only way to share information is by making disks and copying files between the two networks---decidedly low tech!

Separate Networks



The next advancement in Industrial network design was the bridge. Bridges act as a "gatekeeper" between two collision domains. By being physically wired into both LANs, this device is able to discern the source and destination address of an Ethernet packet. The bridge is also capable of "mapping" the locations of Ethernet nodes on either side of itself. By linking a control network and an office network with a bridge, you can stop traffic that is meant to travel between two computers in the office LAN from burdening devices on the other side of the bridge. When traffic occurs that is addressed for a device on the other side of the bridge from the originating address, the bridge will allow this traffic to pass. Compared to the completely shared network, the bridged network can reduce, but not eliminate, the opportunity for collisions and network slowdowns.

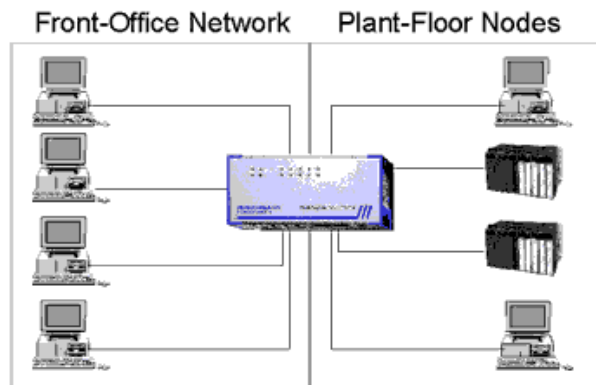
Bridged Networks



The newest generation of networking equipment, however, combines the multi-connectivity of the hub with the selective routing of Ethernet packets of bridges. A switch is generally a multiport device which has the ability to "read" the address portion of an Ethernet packet and then send the packet out the port on which the destination node resides. Think of a switch as a multiport bridge. Most modern switches have buffers that allow them to store and forward the Ethernet packets that are sent to it. Each port of the switch can connect either directly to a node or to a hub(s) which can also have multiple nodes connected to it. Modern switches also have plug-and-play capability. This means that they are capable of learning the unique addresses of devices attached to them (even if those devices are plugged into a hub which in turn is then attached to the switch) without any programming. If a PC or PLC is plugged directly into a switch, the switch would only allow traffic addressed to that device to be sent down the connection cable to the device. By controlling the flow of information between ports, switches achieve major advantages over current shared environments:

- When all devices are directly connected into a switch port, the opportunity for collision between ports is eliminated. This assures that packets will arrive with much greater certainty than in a shared environment.
- Each port has more bandwidth available to it at any time. In a shared environment, any port in the system could consume the entire bandwidth in the network at any point in time. This means that during a peak in traffic, the network availability of any other node is greatly reduced. In a completely port-switched environment, however, the only traffic flowing down the wire between any node and the switch is either traffic destined for, or created by, that particular node.

Port Switched Network



Switches provide Industrial users with many of the safeguards that could only be provided by wiring distinct, proprietary-control networks in the past. The elimination of collisions by connecting every node to a switched port, coupled with the ability to keep control and office traffic from interacting unwontedly, while still using one physical network, allows industrial users to enjoy the open architecture and massive bandwidth and speed of Ethernet without compromising the integrity of their control traffic.