The ABCs of Selecting the Proper Industrial Ethernet Switch

Introduction

Selecting the proper Industrial Ethernet switch can be a confusing task. There are many options to consider such as auto-negotiation features, managed versus unmanaged, redundancy, environmental concerns, future-proofing, determinism issues, and many more.

What Does a Switch Do?

Before investigating all the features provided by Industrial Ethernet switches, we should first consider exactly what an Ethernet switch is and what it does. An Ethernet switch (also known as a switching hub), basically, interconnects Ethernet devices. It receives frames transmitted by one device and passes these frames onto appropriate switch ports which connect to other Ethernet devices. As it passes these frames it also learns where Ethernet devices are located and uses this information to help decide which ports to use for passing frames. This helps cut down on network utilization as frames only go to the appropriate switch ports.

An Ethernet hub (also known as a repeating hub) is a similar device, but it only allows one device at a time to communicate on the entire network. This provides less network efficiency as each device must wait its turn to transmit a frame.

Most Ethernet switches use the store-and-forward method to pass frames. The switch receives an entire frame and then transmits this frame out the appropriate port(s). A repeating hub, however, will receive one bit and re-transmit this bit.

As the switch learns where devices are located, it stores this address information in its built-in address table memory. It also has memory for frames it is in the process of forwarding. Each switch will vary in address storage space and frame storage space. One consideration is how many Ethernet devices you will have in your entire network. Your switch should have enough address storage memory for each device in the network. For example, if you plan on having 100 Ethernet devices in your network, your switch should have an address table large enough to store at least 100 Ethernet addresses.

Most switches provide basic functionality for monitoring the network in the way of LEDs for link and activity. Most switches also provide the ability to negotiate some of their settings.

Let's Negotiate

Many switches support two types of negotiation. The first is called "auto-negotiation." This protocol allows two interconnected Ethernet entities (for example a switch and a computer) to come to some agreement over several parameters used in their communications. One parameter they negotiate is the data rate to be used. This is generally 10 Mbps or 100 Mbps. They negotiate the use of half- or full-duplex. Full-duplex allows communications to exist in both directions at the same time, while half-duplex only allows communications in one direction at a time. They also negotiate the use of flow control in their communications. If flow control is utilized, then either device can request that communications be halted if the device needs time to process received frames.

If supported, Auto-MDIX or Auto-crossover is another negotiation that can occur between Ethernet entities. This allows these entities to decide which wire pair to use for transmitting frames and which wire pair to use for receiving frames. This feature is attractive when connecting two switches as this would normally require the use of a special cross-over cable. With Auto-MDIX the two switches negotiate which wire pairs to use when communicating and this allows the use of standard (straight-through) cables when connecting two switches or two end devices.

Fiber or Twisted-Pair

Most Ethernet communications occur over twisted-pair cabling. However, there are times when communications should occur over fiber-optic cabling. These cables are used when signals need to travel over greater distances than the 100 m supported by twisted-pair cabling. Most fiber-optic devices can communicate up to 15 km when using a full-duplex setting. Fiber-optic cables are also used in high noise environments because their communications are unaffected by electrical or magnetic fields.

Managed vs. Unmanaged

One of the key questions in choosing an Industrial Ethernet switch is whether to select a managed switch or an unmanaged switch. A managed switch is generally more expensive than an unmanaged switch. However, with this additional cost, extra features are provided. Also, a managed switch will contain the features of its unmanaged siblings.

A managed switch is basically a switch which supports the SNMP (Simple Network Management Protocol). Of course, most managed switches provide features beyond SNMP.

Basically, a managed switch allows you to take control of your network. An unmanaged switch will simply allow Ethernet

devices to communicate. You connect your Ethernet devices to the unmanaged switch and they usually communicate automatically. There will be status LEDs to give you some feedback regarding link and activity, but this is generally all you get. With a managed switch you will have the same status LEDs, but the managed switch will let you adjust your communication parameters to any settings you desire and let you monitor the network behavior in a number of different ways.

For example, in systems that communicate in high noise environments, it is sometimes advantageous to force the data rate to 10 Mbps because noise coupled into the cables may confuse the auto-negotiation process. Most managed switches will allow you to set the data rate of each port. These environments can also benefit from disabling Auto-MDIX support since this negotiation can become confused by noise. Again, a managed switch is normally required if you want to enable or disable this feature on a port-by-port basis.

With a managed switch you can also monitor the network. Through SNMP you can view a multimode of network statistics. This includes the number of bytes transmitted, received; number of frames transmitted, received; number of errors and port status. All of this can be viewed on a port basis. Some managed switches also make this data available via a web server so that you can use a standard browser to view the network status.

Most managed switches also offer advanced features that enhance your control of the network.

Advanced Switch Features

Features such as Quality of Service (QoS), Trunking, Virtual Local Area Network (VLAN), port mirroring, fault relay, IGMP snooping, redundancy and SNMP are normally only found on managed switches.

QoS

QoS is the ability of the switch to apply a higher priority to certain frames. A switch can use the port on which the frame arrived to determine the frame priority (port QoS) or it can use a tag within the frame to determine its priority (IEEE 802.1p and 802.1Q). These features are useful in improving determinism.

Trunking

Trunking is two or more ports grouped together and acting as one logical path. This can be used to increase the bandwidth between two switches. Also, in some cases, these paths can provide some redundancy. For example, if two 100 Mbps switches are interconnected by two cables, then the bandwidth between these two switches can be 200 Mbps.

VLAN

VLANs allow a switch to logically group devices and to isolate traffic between these groups even if the devices all share the same physical switch. For example, if the switch was being used for both office communications and factory communications, two VLANs could be created to isolate the office communications from the factory communications. Some switches also allow devices to be located on multiple VLANs. This is sometimes called overlapping VLANs. This can be used if one device, a SCADA system for example, needs to communicate to both the office and the factory then this device would exist in both the office VLAN and the factory VLAN. This would isolate traffic between the remaining office and factory devices, but allow the SCADA system to communicate on both networks.

Port Mirroring

One of the advantages of using a switch instead of a hub, which is traffic isolation, also becomes a disadvantage when trying to debug a communications problem. A switch will only send frames to those devices in the conversation. This helps lessen network congestion. However, if you want to view all frames transmitted on the network, this feature becomes an issue. Many managed switches offer a feature called port mirroring. Port mirroring allows one port of the switch to monitor the traffic sent/received by one or more ports of the switch. With this feature a PC running a protocol analyzer program can capture traffic from one or many ports after port monitoring has been enabled. Protocol analyzers are popular problem-solving tools for Ethernet networks.

Fault Relay

Many switches offer a fault relay to monitor link status of specific ports or the power status of the switch. These dry contacts can then be connected to a PLC or other such control device and used as an input to the control system. This is useful if you want to alert the control system when communication to one or more Ethernet devices has failed.

IGMP Snooping

In Ethernet networks there are three types of frames—broadcast frames which are destined to all devices in the network, directed frames which are sent to one specific device, and multicast frames which are sent to one or more devices. Some Ethernet protocols utilize multicast frames to send data to multiple devices at the same time. These protocols generally create a large amount of multicast traffic. Switches with IGMP (Internet Group Multicast Protocol) snooping can automatically send the multicast frames only to devices which have requested these frames. This keeps the multicast frames from flooding devices which have not requested these frames. Some devices may be unable to perform their normal activities when they receive a large amount of unwanted multicast traffic. This multicast filtering can be important in large Ethernet/IP networks.

Redundancy

Redundancy is a popular feature in managed Industrial Ethernet switches. Basically, they provide the ability to interconnect these switches in a manner such that if one interconnecting cable were to fail, another cable or set of cables would take over. The time in which this recovery takes place is called the recovery time. There are two popular IEEE redundancy standards: Spanning Tree Protocol (STP) and Rapid Spanning Tree Protocol (RSTP). STP (IEEE 802.1D) is the older, slower-to-recover protocol. RSTP (IEEE 802.1w) is the newer, faster-to-recover version of STP. STP generally can recover from a fault within 30 to 60 seconds. RSTP, generally, can recover in 1 to 2 seconds.

Because of these long recovery periods, many industrial switch vendors have created their own proprietary ring redundancy protocols. These can generally recover in less than 300 ms. In these networks, the switches are connected in a ring (see figure below). These protocols generally select one switch-to-switch link in the network to be disabled. This is the backup link. When another switch-to-switch link fails the backup link is enabled, thus repairing the network.

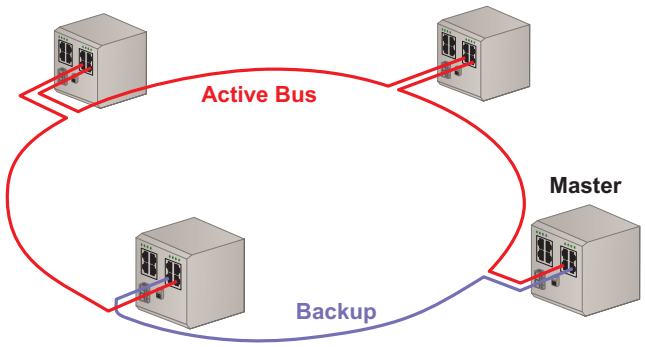

STP/RSTP vs. Proprietary Ring

Each redundancy method has its advantages and its disadvantages. STP and RSTP can be wired in a ring or in just about any configuration imaginable. There are several vendors selling STP and RSTP-compatible switches. STP/RSTP networks, however, are generally slower to recover then proprietary ring networks. Proprietary Ring protocols must be wired in a ring or in several rings and all switches in the ring must come from the same vendor. Also STP/RSTP networks can provide faster recovery times if you utilize a mesh network. This requires three connections between switches, while the ring network only requires two. You should consult your switch vendor for their network recovery time.

SNMP

SNMP is the Simple Network Management Protocol that allows network management applications to communicate with a switch in a standardized way and request status information and set configurations. Most switches also support a similar set of management data (MIBs). As most managed switches support this protocol and support similar data, one application can communicate with multiple switches. This allows one application to monitor multiple switches at the same time, thus providing a global view of the network. Network management applications provide many features such as the ability to graph data collected from one or many switches, the ability to map the network, the ability to receive error frames from the switches, etc. There are also a few vendors who sell applications which convert SNMP data into OPC data for use in HMI systems.

Protocol Required Features

In some cases industrial equipment vendors or industrial protocols require features only available on managed switches. For example, several EtherNet/IP control vendors recommend the use of switches that support IGMP snooping. Also, your vendor may suggest that your switch supports QoS or VLAN. Consult your equipment vendor for any suggested or required switch features.

Environmental Concerns

This is where Industrial Ethernet switches really differentiate themselves from commercial switches. Industrial Ethernet switches were designed for environments that are not favorable to commercial switches. This can include environments with temperature extremes, high vibration and severe electrical noise. As commercial environments are generally room temperature, most commercial switches are designed for a very small temperature range. Also, some commercial switches use fans to help in cooling. This could be a problem in many industrial environments due to dust that could accumulate in the fans and the low MTBF of most fans. Commercial switches also expect to be in a low electrical noise environment. This is not the case in many industrial environments. Commercial switches are designed to meet commercial or office EMC requirements while most Industrial Ethernet switches are designed to meet the more stringent industrial EMC requirements.

Environmental concerns can also be as simple as mounting and power supply issues. Most commercial switches are designed for 19-inch rack mounting or table top mounting, but this usually is not acceptable for many industrial environments. Normally these environments utilize DIN-rail mounting or panel mounting. Also, many commercial switches utilize a wall-mount power supply. These are generally hard to install in industrial settings, and they can be problematic in high-vibration environments where vibrations can dislodge them from standard electrical outlets.

Future-Proofing

How do you future-proof a system? Let's say, for example, a small system was put together using an unmanaged switch. Initially, this may be an acceptable situation. However, as more Ethernet devices are added, the system becomes more complex and the need for a managed switch may become more apparent. Having a managed switch in place initially may help future-proof your system so as more devices are added, you will retain the ability to fully control and monitor your network. Also, as you add Ethernet devices you may find that QoS can be useful to help prioritize the frames due to the increased traffic load of your network. You may also want to utilize VLANs to help isolate devices due to the increased number of Ethernet devices on your network.

In the future, Ethernet protocols may require QoS standards such as 802.1p/802.1Q to help achieve higher priority for frames traveling on the network. VLANs can be used to isolate devices which may be sensitive to the higher levels of traffic as the network grows in size. Having a switch with these features now helps to future-proof your system.

If you are considering moving to EtherNet/IP networks in the future, it might make sense to use a switch that supports IGMP snooping. It is possible that the initial system is small and isolated from other networks and does not require the use of IGMP snooping, however, later it may become desirable to interconnect this system to the company network. At this point, without IGMP snooping, you may have a large amount of

multicast traffic being sent to the company network. A switch with IGMP snooping could help to block this traffic. Also, a switch with VLAN support could be utilized to block this traffic from entering the company network.

Determinism Issues

Ethernet switches can support a deterministic system. The question you need to ask is what type of response and jitter is acceptable in your system. If I send a frame from one device to another, what is the maximum response time and what is the maximum acceptable jitter? Most Ethernet switches utilize store-and-forward when processing a frame. This means that the entire frame is received and then re-transmitted. If you cascade several switches, the delay in storing-and-forwarding these frames increases for every switch. Also, each switch has a small amount of internal latency which adds to this delay.

However, if your system is communicating at 100 Mbps, then the time for the store-and-forward of each frame is generally on the order of 10 ms for small frames and 130 ms for maximum sized Ethernet frames. Also, the jitter is fairly small, approximately 1 ms per switch. This delay is generally very small compared to the latency of most TCP/IP devices. And protocols such as Ethernet/IP are adding IEEE 1588 (Precision Clock Synchronization Protocol for Networked Measurement and Control Systems) to help with synchronizing Ethernet connected devices.

What Should I Choose?

As you can see, there are many factors to consider when purchasing an Industrial Ethernet switch. We have discussed a number of features found on many Industrial Ethernet switches. Please consult your vendors to determine the features supported on their Industrial Ethernet switches. As there is no standard that dictates all the mandatory features provided on an unmanaged or managed switch, you will need to ask your switch vendor which features they include in their products.