

Switch with Confidence: Avoiding the Pitfalls with Industrial Managed Switches

Written by: Paul Wacker, Advantech Corporation, Industrial Automation Group

The office isn't a factory but its networking standard, Ethernet, can be successfully applied in a plant – if the proper steps, and some care, are taken. Done right, the benefits include an inexpensive communication and control network with a low total cost of ownership and virtually 100 percent uptime. Done wrong, the results could be unexpected downtime, a loss of product, and the kind of events that make the news.

A key to getting things right is the use of suitable technology, such as industrial managed switches. These devices are ruggedized and more intelligent than their unmanaged cousins, a difference that enables them to offer advanced traffic control, greater redundancy, and better security, as well as remote monitoring and management. All of those characteristics are important to successful networking and can help avoid pitfalls.

This white paper looks at how industrial managed switches, together with the right network topology and an understanding of failure modes, can lead to almost 100 percent network uptime. First, though, the basics of industrial Ethernet will be explored. The paper will wrap up with an examination of security concerns and tips for new projects, such as designing in redundancy of power supplies and media as steps to eliminate downtime.

Not your father's Ethernet

Before showing how managed switches can be used to advantage in an automation setting, it helps to look at the differences between industrial and traditional Ethernet. In general, industrial Ethernet technology is ruggedized, for use in harsh environments, has special packaging and longer product life cycles than its office-targeted counterparts.

Why use a version of a networking technology not originally developed for industrial applications? That's because Ethernet:

- Is inexpensive compared to proprietary communication and control networks.

- Is well established, based on open standards, and has good market acceptance.
- Allows multiple protocols to exist on the same wiring and enables easy integration of systems and components from multiple vendors.
- Provides future-proof compatibility, very flexible network topologies, and is highly scalable.

The result is a low total cost of ownership, which makes Ethernet an attractive network solution. However, it was originally designed for an office setting, not a factory floor where dust, vibration, and temperature extremes are the norm. This explains the need for adaptations to accommodate plant and factory environments.

Industrial Ethernet looks exactly like the office variety, at least with regard to signals on the wiring. As for the wiring itself, the two can also look quite similar. Standard Ethernet Category 5 (Cat5) and later (Cat5e and Cat6) cabling uses twisted-pair wiring that is transformer coupled, which provides noise immunity and ground fault protection required by industrial networks.

In other circumstances, though, industrial Ethernet demands a different wiring solution. For runs longer than 300 feet, or about 100 meters, optical fiber is the best choice in an industrial setting, as it offers total electrical isolation in electrically noisy environments. Optical fiber may also be appropriate for shorter runs, if the situation warrants it.

Because of vibration, solid wiring should never be used on a factory floor. For the same reason, special attention must be paid to cable terminations and connectors. For the former, a quality terminated cable or one adhering to the RJ45 standard should be used. For the latter, bulkhead connectors should be used beyond a panel. When it comes to machine-mount devices, specialized connectors, like RJ45 with a screw-on cap or the M12 industrial connector, should be used.




Figure 1 Connectors: RJ45 (top) and M12 (bottom)

As might be expected, industrial Ethernet has the form and fit required for a factory. This means that components are DIN-rail or panel mounted and has a lower voltage power input, such

as 24 volts DC.



Figure 2 DIN-rail Mount

They are also ruggedized, which means their protection against electromagnetic and radio frequency interference (EMI and RFI) is enhanced. Because plant floors can be harsh environments in terms of electrostatics, part of the ruggedization extends to improved protection against electrostatic discharge, the familiar zap caused during the winter by scuffing feet across a carpet in dry air.

Similarly, factories can be challenging with regard to temperature and vibration, so industrial Ethernet components must function in these extremes. This is accomplished by shock and vibration protection and an extended temperature rating.

The latter has to be accomplished using convection cooling wherever possible because fans are a possible failure point and thus need to be avoided. This is the basis of the requirement for passive cooling using convection. The drive for reliability leads to the use of such things as redundant power inputs and media that can recover from or heal a problem.

A final difference between industrial and traditional Ethernet lies in the life-cycle of products. Whereas those intended for the office may be obsolete in a few months, the availability of industrial Ethernet products is measured in years, with a 5 to 7 year life-cycle being typical.

Finding the best network topology

The shared features of industrial and traditional Ethernet extend to the network topologies. The best choice, like beauty, lies in the eye of the beholder. Thus, there is no “right” network topology.

Fortunately, Ethernet is extremely flexible, which means there are a lot of possible choices. At least one, and probably more, will almost certainly be the right one for a given situation.

Ethernet connectivity can be by traditional wiring. In that case, the warnings and best practices with regard to cable installation and terminal connections apply.

No matter the connection, there are two key characteristics of industrial networking that need to be accounted for. The first is that almost all traffic is local, with something like 80 percent

of it generated by and intended for nearby devices on the plant or factory floor. The second is that service interruptions aren't acceptable. In a standard office situation, a few minutes of email or web downtime is not generally a significant problem. In industrial setting, data must be delivered on time, every time. If not, product may be ruined or an unsafe condition develops.

As for possible network topologies, those vary by application and include

- Star.
- Tree.
- Mesh.
- Ring.

These configurations can be mixed and matched to give the best possible arrangement.

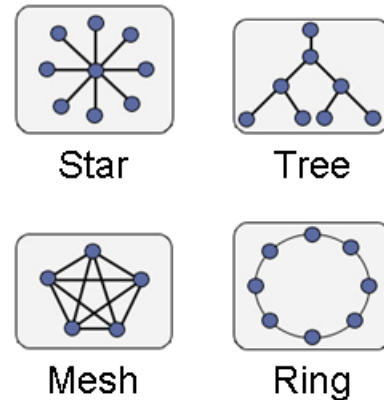


Figure 3 Network Topologies

A star topology works well for simple and small networks. In it, all nodes connect back to a central switch through a single cable. Because of its simplicity, a star layout is the easiest of all to implement. It would be suitable within a single panel or between just a few panels.

A tree topology consists of multiple star configurations connected together. This approach is appropriate for small projects involving two to three switches, where multiple switches are connected together to expanding the number of connected devices or distance between switches.

A mesh star involves multiple paths between some or all nodes on the network. Thus, a device may connect to a switch and to another device. There's a hierarchical structure, with some nodes subordinate to others. This layout may be difficult to implement on a factory floor and the failover mechanisms can be complex. Thus, in the event of a failure of a node, recovery can take seconds, during which time the network may be down and unusable.

As the name implies, in a ring topology all nodes connect to two neighbors, with the result being path that runs through all of them. This topology works well for the overwhelming majority of plant applications. It offers a more consistent performance and fewer bottlenecks than the other

topologies. It also has extremely fast recovery time, taking as few as 10 milliseconds to overcome

single-point failure such as a cable break, disconnection, or loss of power to switch.

Managed vs. Unmanaged Switches

No matter what topology or mix of topologies is used, the type of switch selected can make a crucial difference between a successful network and one that is not. Switches process the basic

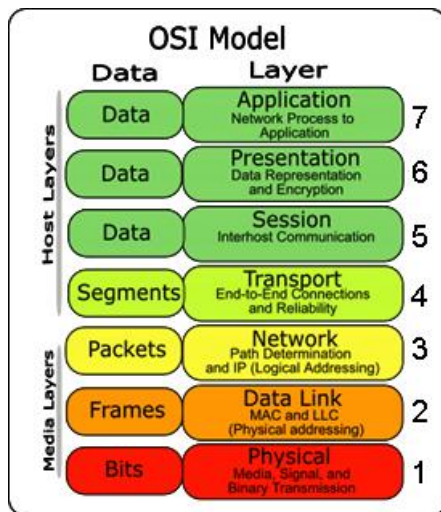


Figure 4 OSI 7-Layer Model

Ethernet communication unit, a frame, at the physical and data-link layers, the first two in the seven layer OSI communication model. Switches are agnostic to higher layer protocols and have two basic transmission modes: one to one unicast or one too many broadcasts. Like all Ethernet devices, switches are identified by their MAC, or media access control, address.

When it comes to switches, the choice is between an unmanaged and a managed one. The first is a plug-and-

play device that acts as network connection point. It receives a

frame from a source and sends it to its destination.

A managed switch, in contrast, is a more intelligent and robust device. It has its own processor and that intelligence brings a number of benefits:

- Manual port settings and advanced traffic control.
- Improved security.
- Remote monitoring and management.
- Greater robustness and media redundancy

Manual port settings mean that port configurations on the switch can be fixed. Unused ports can be disabled for enhanced security. Ports in use can be set for a specific speed or for full or half duplex operation, in situations where devices can't auto negotiate the speed of mode of a connection. Network traffic can be prioritized, with more important messages not having to wait for less important ones to be transmitted and received. Finally, advanced traffic control can make sure

that I/O and similar data does not flood the network. Such information may need to go to multiple

destinations but not to every node, so something is needed that lies between unicast and broadcast operation. Selective broadcast is possible with advanced traffic control, like IGMP Snooping.

As for security, a managed switch can limit access by Ethernet MAC address, using an allowed and blocked list. This method secures the network from unauthorized access. Another means of increased network security is a virtual LAN, a logical network that runs within a single physical network along with other logical networks. This restricts users and devices to only particular parts of the network.




Figure 5 Virtual Local Area Network (VLAN)

A managed switch can mirror traffic from one or more ports to another, enabling remote traffic monitoring. A managed switch itself can be managed remotely, though a web interface or SNMP, the simple network management protocol standard

As for robustness, a managed switch will often have redundant network connections. Thus, if one is lost, the switch can still function. The on-board CPU allows it to rapidly recover from a failure because the processor can quickly implement corrective action. In a ring configuration, for example, if a link goes down than the switch can reach a device on the other end and thus heal the breach. Besides self-healing rings, managed switches can implement the Spanning Tree and Rapid Spanning Tree Protocols (RSTP) for recovery.

Designing for 100% uptime

The benefits of managed switches are not free. A managed switch is more expensive than an unmanaged one.

Against that, though, must be weighed the cost of downtime. That will vary for different situations and different plant floors, but in many cases it will be substantial. That's particularly true for a factory full of expensive equipment, which may sit idle due to a network problem. Because of their capabilities and design, managed switches can play a key role in helping maximize uptime.



Figure 6 Cost of Downtime

Once the full cost of downtime is understood, then the cost effectiveness of attempts to eliminate it can be judged. It is critical in designing for maximum uptime that the key failure modes are accounted for.

The first failure mode happens before the network is even put into place. It's important to use industrial-grade hardware and cables, with these supplied by a vendor with experience in industrial automation.

Even the best hardware and cables can fail, so there's a need for redundancy in those areas most likely to do so. One of these involves power supplies, which are usually the most common point of failure. Thus, having two or more power supplies in a switch can ensure that downtime is minimized.




Figure 7 Self-Healing Ring

Another likely point of failure involves the links between devices. Cables can disconnect or break, which makes media redundancy important. This can be achieved by trunking, which physically duplicates links, or through the use of a self-healing ring using RSTP or a vendor specific implementation for more rapid recover.

What about security?

At one time an afterthought, the security of an automation network is now increasingly important. This greater emphasis is partly due to the success in networking the factory floor. It's ©2009 Advantech Corporation, Industrial Automation Group.

now possible to access devices from around the world, which makes them much more vulnerable to casual hacking or outright attack.

In some ways, security improvement ties into uptime efforts, since a compromised network is more likely to go down than one that has not been infiltrated. Here, again, managed switches can play an important role.

Because they allow for the use of VLANs, managed switches can ensure that maintenance personnel or the HMI for operators on the line can only get to those devices that are needed to do the job. By disabling unused ports, managed switches can decrease the area vulnerable to attack, making an assault harder to carry out.

Nonetheless, security is a challenge, in part because many automation products and devices on the factory floor were not designed with security in mind. Such basics as authentication and authorization protocols may be weak or altogether missing.

For that reason, some additional guidelines to follow to enhance security are:

- Establish a security policy, with awareness and training.
- Run periodic self-assessments.
- Restrict physical access to switches and cabinets.
- Provide dedicated ports for local work by OEMs.
- Collaborate with IT.

Conclusion

The last tip is a particularly useful one, as it can help in the deployment of new networks. Getting to know the IT staff can help in aligning projects to business objectives and can allow for some two-way learning.

A second tip is that it's important to design for future expansion. This means that there should be plans for extra ports to account for additions to a panel, the inclusion of new panels, or the incorporation of laptops and workstations.

A third tip is to remember that downtime is expensive. Thus, it makes sense to design for high availability. Redundant power supplies and media are two important ways to do so – but not

the only ones.

Industrial managed switches can help ensure that downtime and other problems are avoided. Thanks to their on-board processor, these devices can provide advanced traffic control, greater redundancy, and better security, as well as remote monitoring and management. By taking advantage of these capabilities, you can avoid networking pitfalls.

###