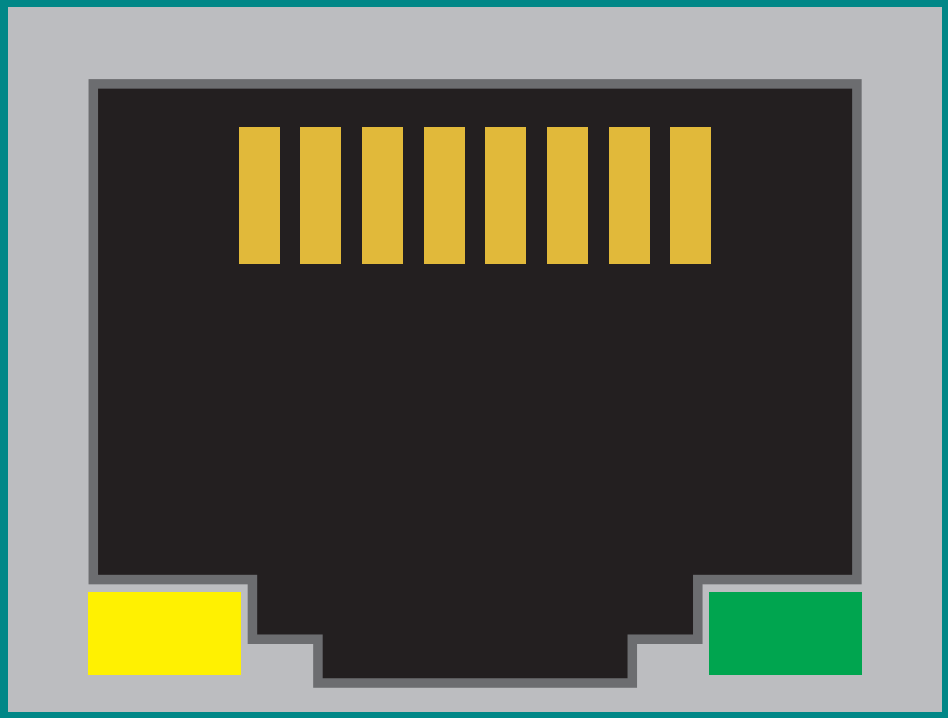

The Engineer's Guide to Industrial Networking

Best Practices to Define & Design a Reliable Industrial Network

MOXA®

The Engineer's Guide to Industrial Networking

Best Practices to Define & Design a Reliable Industrial Network

Is This E-book Right For You?

The objective of this guide is to aid in the proper design of an industrial network. Reading through the document will arm the reader with the knowledge needed to architect an industrial network. Industrial networks have a different set of requirements than what is found in enterprise. In the enterprise space, network failures typically have marginal impact to business operations - think of losing emails. In an industrial network, failure can be extremely expensive, and can even result in physical injuries in extreme circumstances.

The first part of this guide covers best practices when beginning to design a reliable industrial network. We will explore defining network requirements as the starting point with a focus on identifying the needs and wants of a network. A methodical approach to achieve this requirement is explained, and in doing so, QoS (Quality of Service), costs of a network, and general requirements will be discussed.

The second part of this document covers best practices that include: when to choose wireline vs. wireless, network redundancy, segmentation, multicasting, ease of management, and security. Topics are explored in technical detail with an emphasis on industry specifications.

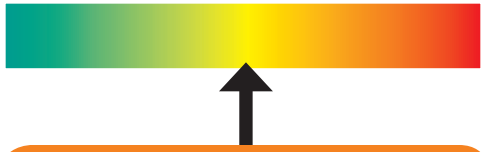

This guide is rated intermediate. It is recommended for those who have basic knowledge of networking technology and are looking to build/expand their industrial network.

After reading this guide you should be able to:

- **Avoid common pitfalls in network design**
- **Begin to architect a new or expanding network**
- **Understand the unique requirements of industrial networks**
- **Understand basic networking concepts**

Table of Contents

Introduction

Networks serve to bridge all communication within an industrial system. For data to get from one point to another, it will most likely have to pass through a network. The network serves as the data highway and digital nervous system by interconnecting intelligent devices and providing important data to all network nodes, whether it is a PLC, HMI, or a PC. There have been great advances in the last couple of years in the industrial space. Not long ago, industrial communication meant a field protocol using RS-422 or RS-485 serial connectivity, however this is evolving.

Networks have progressed and become more deterministic and reliable, and this has pushed the industrial space to use networking infrastructure. There have been great advances as communication systems are now utilizing standards-based communication that can serve an industrial network and more. The aggregation of separate communication systems has led to the converged network, a network that performs multiple duties from servicing SCADA, voice, video, and more. With this convergence, networks are often treated as any commercial network, with less than optimal results.

The industrial space requires far higher standards than what is common in the enterprise/commercial space. For enterprise applications, typical network outages result in undeliverable emails and loss of web access. It is usually acceptable to have a few minutes of unpredicted network downtime in a commercial environment. For the industrial space, this is a different story; networks need to be reliable and deterministic. For the automation industry, unreliable networks can cause more than just financial loss, although a loss of communication for just one minute can already have strong financial impacts. An unpredictable network is unacceptable in automation, where determinism is paramount. With industrial networks, failure can be extremely expensive, and even result in the loss of life. In short, failure is not an option for mission-critical networks.

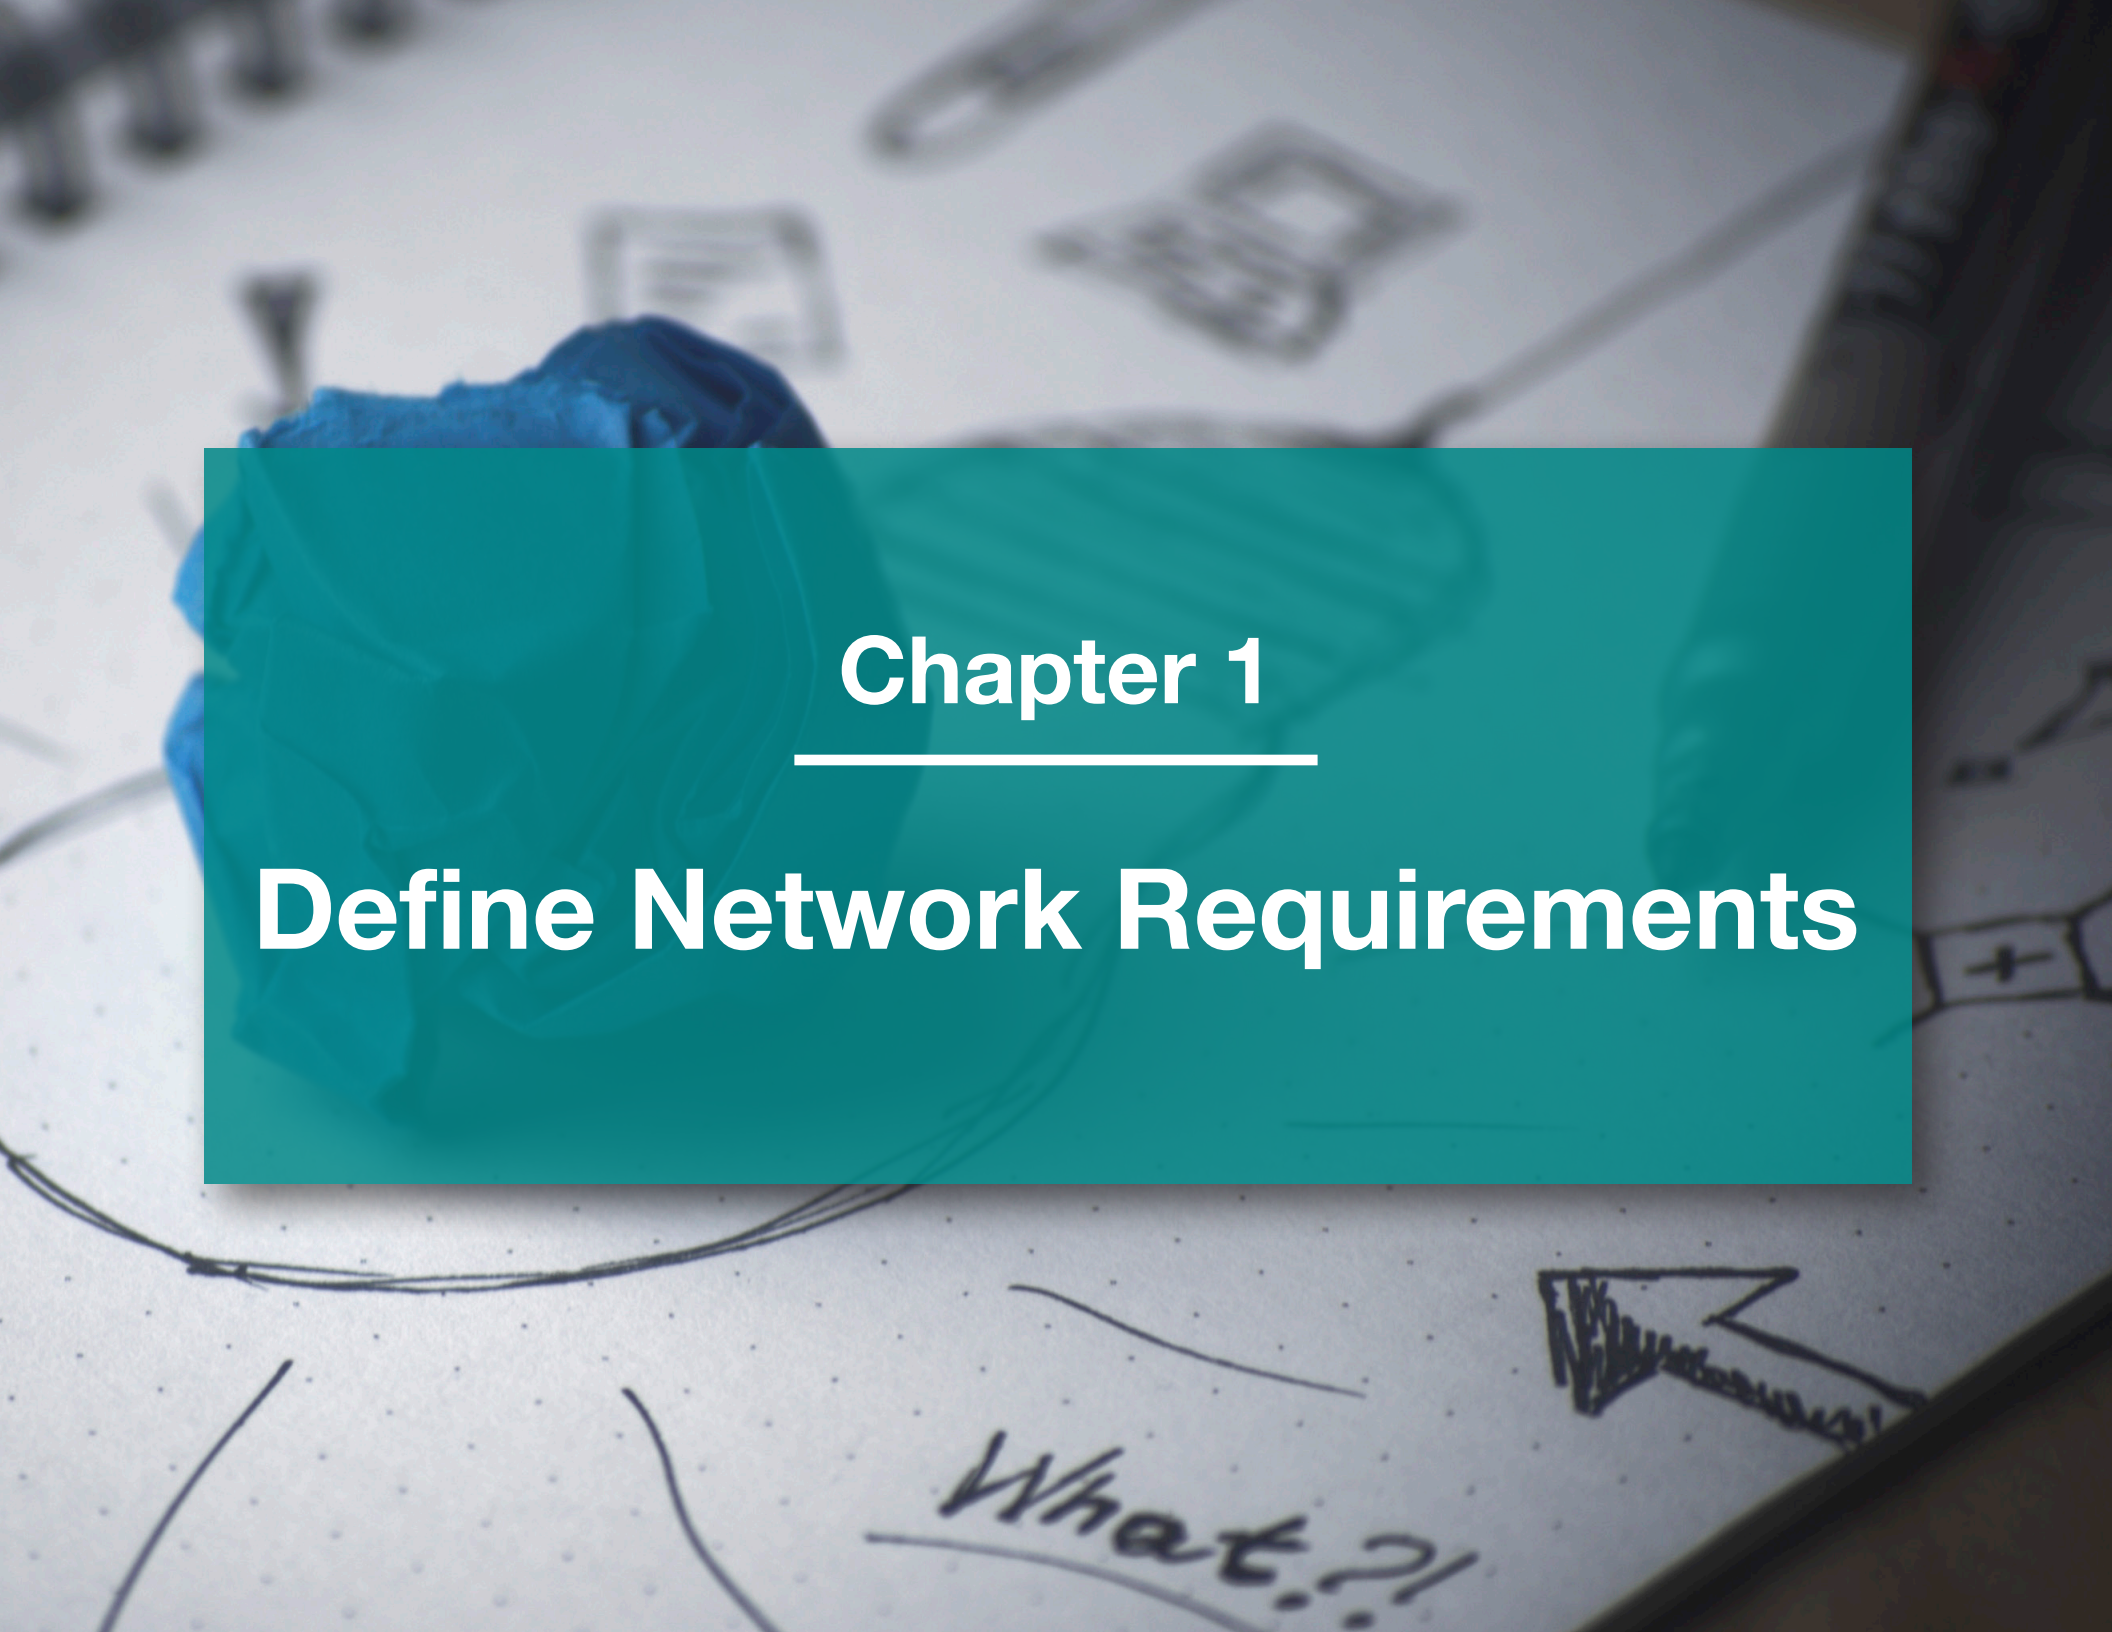

Chapter 1

Define Network Requirements

Define Network Requirements

Before setting out to build a network, one must understand the network requirements. This means figuring out what is going to be connected to the network and the intended purpose of the network. Requirements change based on what is connected, and addressing the different challenges is core to a successful network deployment. The intended role a network serves is also critical. For example, if it's in a mission-critical environment, then network uptime is of even greater importance.

Remember that industrial-class networks have very different requirements than those of the commercial/enterprise space. In the commercial space, network failure and loss of traffic can be tolerated to an extent. Should a catastrophic network failure occur, operations can be vulnerable, but no one gets injured. In an industrial network, people can get hurt or worse.

At the very core, the requirement for the system including network size, application requirements, cost of downtime, tolerance for data loss, and level of determinism should be analyzed. One of the first and most important requirements for design consideration is QoS.

QoS Defined

The recent upsurge of industrial networks has added to the increased importance for Quality of Service (QoS). QoS is defined as the overall performance of a network, above all the perceived performance seen by the users. The keywords here are "perceived performance". Also, notice that protocols are nowhere to be found in this description. Despite popular belief, Quality of Service is not a protocol stack or protocol of any kind, it is a methodology. That said, there are protocols that help ensure QoS, but that is a topic for later discussion.

QoS is a quality parameter and is the categorical evaluation of network traffic. It is important to note that in a network, QoS is only affected when there is network congestion. Network data in a router or switch is handled in a FIFO (First In, First Out) fashion, where incoming data is buffered and leaves in the order it was buffered. If there is more incoming data than the outgoing bandwidth can support, the network device is forced to buffer the data and network congestion occurs. When there is network congestion, QoS becomes relevant and works to differentiate traffic from best-effort, which means it lacks any guarantee that data will get there in a timely manner. Certain traffic streams are either prioritized or have reserved bandwidth, meaning the devices that are tagged will have guaranteed bandwidth allocated to them.

Some network congestion is expected. The real problem arises when there is too much congestion. If the buffer overflows, the network device will begin to drop packets. For this reason, the Internet, as well as network traffic is traditionally known as “best effort”. Even if packets are not dropped, if they take too long to get to their destination, packets may have been retransmitted, which adds to congestion.

> Connection-based protocols such as TCP have a timer that waits for an acknowledgment from the receiver. If the receiver does not receive the packet before the timer expires or if the acknowledgment packet itself is delayed or dropped in congestion, the result is that the sender will retransmit the packet (it will in fact retransmit a set of packets at a reduced window size).

When talking about QoS, it is best to talk about determinism first as it is an important trait in many industrial networks. Wikipedia defines a deterministic system as: “a system in which no randomness is involved in the development of future states of the system”. Thus, a deterministic system is one with predictable behavior. A network that is deterministic has behavior that can be expected and therefore designed around. In order to predict network behavior, a set of metrics need to be established. For QoS there are three primary metrics: delay, jitter, and reliability.

Delay

In networking, delay refers to the amount of time it takes for the data to get from point A to point B, which can incorporate a great deal of things. This is most analogous to the time it takes you to commute to work. On a good day, it may take you 30 minutes. In networking, data transverses a network much faster, nonetheless any delay can radically change the network behavior.

Jitter

Jitter is the variance of delay. Using the same analogy of your daily commute, now factor in that traffic can be unpredictable at times. Normally, it should take you 30 minutes, whereas it may take you 20 minutes on good days or 60 minutes on bad days. This variance or inconsistency is referred to as jitter. If delay was a fixed time interval, this means it is predictable and therefore deterministic. Jitter makes the system less predictable. The more jitter, the more difficult it is to build a deterministic system. In video and voice communication, buffers are introduced to cope with jitter, however, this increases the delay. Similar methods can be applied to communication in general.

Reliability

The third metric for QoS is reliability. The reliability parameter is also referred to as the “Loss Ratio”. This parameter defines the average error rate. An example of an error that can occur is the packet or frame not arriving at its destination or not arriving in a timely manner.

QoS Approaches

The recent upsurge of industrial networks has added to the increased importance for Quality of Service (QoS). QoS is defined as the overall performance of a network, above all the perceived performance seen by the users. The keywords here are “perceived performance”. Also, notice that protocols are nowhere to be found in this description. Despite popular belief, Quality of Service is not a protocol stack or protocol of any kind, it is a methodology. That said, there are protocols that help ensure QoS, but that is a topic for later discussion.

There are three methodologies used in order to achieve set quality parameters: Predictive Quality, Flow-based Quality, and Cloud-based Quality.

Predictive Quality

With Predictive Quality, a network is designed so that it is over-provisioned. Over-provisioning a network requires designing it so that congestion will not occur. This can be done by using faster interconnection and port-aggregation where necessary to scale up. This is synonymous to the lines one experiences at an airport. In an airport there are usually multiple lines, whether it is at check-in or before boarding. Most airlines have multiple lines for the boarding process. You can usually pay more to be upgraded to a higher-priority line. This is not unlike many amusement parks and tourist attractions.

Imagine if you can see the line before making the decision to pay extra. If there is no one in either line, then you know you will go in quickly. That is what over-provisioning is about - making sure that lines or queues as it is known in networking do not happen. This is done by ensuring that the pipe is much bigger than the load. An alternative example can be a freeway in which there are more lanes than necessary. There is a downside to this approach, that is as a network's demand increases, so must its bandwidth. This means equipment may also need to be upgraded to keep up with the requirements of the changing environment. Fortunately, industrial networks seldom change. For this reason, Predictive Quality is the most widely used form of QoS in industrial networks.

Flow-based Quality

A stream of packets or frames with the same quality requirements is a flow. In Flow-based Quality, network resources are provisioned or reserved during each connection. This requires a signaling protocol to establish and breakdown connections. The advantage of this is that when fully supported, it works reliably. By creating dedicated connections, it forces a packet switching network to behave similar to a circuit switching network.

The most used Flow-based Quality protocol is "Integrated Services" or IntServ for short. IntServ is readily found in many networking devices. It relies on the RSVP (Resource ReSerVation Protocol) as the signaling protocol used in establishing connections and resource reservations. The problem with IntServ is that it only works well in a small scale as it quickly becomes difficult to keep track of all the reservations.

Flow-based Quality does suffer from certain disadvantages. If a flow was long-lived, this approach may work, however this is seldom the case. On the Internet, a flow may contain anywhere from ten to twenty packets, yielding a lot of overhead per flow. In addition, each switch must maintain the state of tens of thousands of connections simultaneously, which also makes this impractical and difficult to achieve. Flow-based Quality is seldom seen in the industrial space.

Class-based Quality

Class-based Quality also relies on protocols like Flow-based Quality, but takes a different approach. Packets are clustered based on similar quality requirements. Within these clusters, it is best-effort. Each cluster is prioritized differently and therefore gets a different level of access to network resources. This methodology relies on the traffic being marked. Based on these markings, the network device chooses how to prioritize traffic. Higher priority traffic is placed into a queue that will more readily send out traffic. Lower priority traffic will be placed into a queue where it will be forced to wait for the higher priority traffic.

To best explain Class-based Quality, let's revisit the airport line scenario. In this example, there are multiple lines for flight check-in based on priority. The ticket tags the passenger to which priority he goes in. It is up to the boarding agent to check the tagging and ensure which queue/line boards first. Within each line, it is best-effort, in which if you were first to the high priority queue, you will leave that queue first. In the example earlier, we talked about over-provisioning. With over-provisioning, there is no prioritization, and instead relies on ensuring no lines. This is synonymous with there being several boarding agents, so many in fact that lines do not form. Should a line form, that is when Class-based Quality takes over. Thus, over-provisioning and Class-based priorities can be used concurrently.

Class-based Quality methods solve a lot of the problems regarding scaling that Flow-based Quality prioritization method suffers, and is argued as being less expensive to implement than the Predictive Quality methodology.

There are several widely used protocols for Class-based Quality. Type of Service (ToS), Differential Services (DiffServ) and Class of Service (CoS) are three of the most recognized. Each of these protocols use markings/tags to determine priority. For more details on the protocols of Class-based Quality, download our white paper: **Understanding the Protocols of Class-Based QoS**.

Total Cost of Ownership

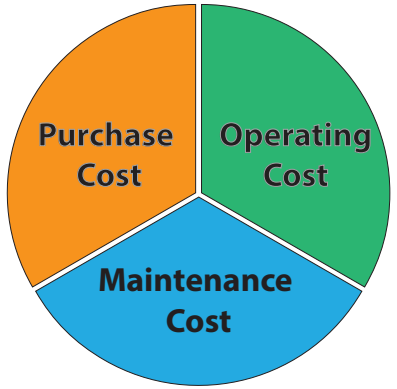

Total Cost of Ownership is often an overlooked yet important step in making a network design decision. It is easy to focus on the acquisition cost or upfront cost and neglect the true cost of the system. Although acquisition cost is important, you will find that cost of ownership can amount to far more than expected. Most costs can be broken down to three categories: purchase, operating, and maintenance costs. All three are important but every application will have a different emphasis on each of these costs.

When choosing equipment, think of whether it will be reliable in the installed environment. Choosing to install commercial grade in an air-conditioned, non-mission-critical, low EMI setting shouldn't be an issue. However, doing so in an environment that is as harsh as those found in the industrial space is sure to fail sooner than desired.

The cost of maintaining the system is also an important consideration. This means that the skill level of the organization and those tasked with maintaining the network needs to be assessed. Training might be necessary to raise technical knowledge of staff members. Although there may be an upfront cost for training, it will result in lower long term costs. Training allows the staff to better operate and troubleshoot issues that occur. It is easy to create a complex network that is fast and reliable, but with the increased complexity, proper knowledge for how to maintain the system is required. This is not a statement for "KISS" (Keep It Simple Stupid) nor is it advocating this ideology, instead it is stating that simplicity or complexity must be assessed when determining the best network. A network might be complex in order to guarantee redundant and efficient communication, or it might be optimized for maintenance costs. In other cases, a balancing act takes place and the network meets both needs.

White Paper: **6 Tips on How to Lower the Total Cost of Ownership of Industrial IoT Networks**

"Cost of downtime" is an important consideration and part of the total cost of ownership. What is the cost of downtime? This refers to the accumulated costs as a result of network outages or catastrophic network failures. These costs are associated with the loss of productivity, labor costs, startup costs, and so on.

Loss of production is the expense of not being able to create revenue because of a network mishap. This is also known as an opportunity cost, which is associated to the forgone value of an alternate. For example, if you hire someone to mow the lawn so that you can watch a sports game, the cost of watching the sports game is how much was paid to the person mowing your lawn.

To illustrate, let's use the case of a manufacturing plant that generates $1M of revenue/hour producing widgets. If each widget costs $100 and you can create 10,000 widgets/hour then your revenue is $1M/hour from production. That translates to $16,666.67/minute, $277.78/second, or $0.28/millisecond. If a network is designed so that it can recover in 30 seconds, it is going to cost the organization $277.78 x 30 = $8,333.40 each time the network is interrupted. For comparison sake, a network that can recover in 20 milliseconds will cost $0.28 x 20 = $5.60 per occurrence. Keep in mind that these are actually conservative numbers. In many facilities, realistic costs can range from $10,000/minute and reach up to $30,000/minute.

Cost of labor is another consideration that can be associated with operation and maintenance costs. When there is a network issue, there is also a labor cost to find the resolution. A common point that comes up is if you already have labor at a facility, resolving network issues should be the responsibility of the on-site staff. Therefore, there shouldn't be additional labor costs incurred. This can be a true statement assuming the labor force is salaried and was just sitting around idle before the occurrence - which of course is highly unlikely.

The reality is that there is an opportunity cost for taking a trained engineer from doing one task to do another. For example, if the engineer is building out a solution that will increase efficiency by $10,000 an hour and has to shift gears to resolve a network issue, this means that the organization also paid an opportunity cost of $10,000 an hour. Another cost that is often overlooked is the stress and fatigue put on employees. The more difficult the problem, and the more critical, the more stress and fatigue that it is going to put on an employee. This does not make for a happy employee, and therefore can affect productivity, or worse, it can affect employee retention.

Loss of customer confidence is yet another cost. If a network issue hinders an organization from delivering goods/services within their customer's expectations, the organization loses credibility, which can lead to damaged relationships. Needless to say, the worst case scenario would be customers leaving to find new vendors for their goods/services.

Size and Scope

It is important to know what devices are on the network, where, and how many. This will help when identifying proper port count and type of switch needed. Also, be sure to design for the future by considering any expansions down the road. If you are designing a larger network, then Layer 3 switches and VLANs should be a topic for discussion. This is also an important time to understand the bandwidth utilization and determine if over-provisioning should be used.

Measuring bandwidth utilization requires knowing what is on the network, and the load the device will put on the network. This can be challenging at times, and may require a best guess based on previous experience. Sometimes, an estimate can be achieved by identifying the load put on the network with one device, and multiplying it with the number of like devices on the network.

Cost of a Security Breach

There is no doubt that a security breach is a bad thing. Putting a dollar figure on the estimated cost of a security breach, which can include downtime, will help in determining how much to invest in making the network secure.

A security breach can be far more sinister than disruption of the network. A key concern in this day and age is protecting intellectual property (IP). The cost of losing intellectual property can be devastating to an organization. Hackers have made headlines breaking into many companies with a sole purpose of stealing intellectual property. This information can be sold to competitors or even used to create a company in direct competition with the owner of the intellectual property.

Leaving your car unlocked on the street with the windows down is an invitation for thieves. There are thousands of corporate and industrial systems out there offering the same incentive to hackers. Locking your car doors provide at least some security, and would be equivalent to a company implementing a firewall. Locking your doors, rolling up your windows, and installing a car alarm would obviously make it much less likely that the vehicle or its contents will be stolen. This is equivalent to a company using the right hardware at every level and the right security policy. Unfortunately, there are not many industrial control systems that actually implement advanced network security measures. This is surprising considering security is the biggest concern for distributed applications. As the number of connected devices increases, so does the number of external users. To enhance the strength of your industrial cyber security, it is best practice to design your network in accordance with ISA99/IEC-62443 standards and recent ICS-Cert alerts.

ICS-CERT
INDUSTRIAL CONTROL SYSTEMS CYBER EMERGENCY RESPONSE TEAM

Network Diagram

The first and foremost task should be to create a logical network diagram. Shown here is a simplified network diagram that is relatively high-level. It is meant to be easy to understand and should provide a quick check to the network design. Usually these network diagrams will include port information while leaving out details you would typically find in a CAD based diagram. As a result, you should have a simplified view of the edge of the network. A logical diagram is not necessarily designed to show the physical location of the network, but how it is "logically" connected. In a way, the logical network diagram is the cliff notes, or the executive summary of a network. It summarizes the network in a simple diagram that can be shared with others quite easily. The logical network diagram is a first stepping stone to diagramming a network. When finished with the logical diagram, then the CAD based network diagram that includes wiring paths, patch panels, device locations, and all edge devices can be constructed.

On the right, is an example of perhaps an overly simplistic logical diagram, but it is a great starting point. This logical diagram does not contain information about edge devices or port assignments, but it is easy to understand. It shows the primary path of communication via solid black and blue lines, with it's backup path represented by the dotted red lines. Drawing out the diagram helps to ensure that the topology is feasible within the technology of choice. More details about backup redundancy technology will be discussed in Chapter 2: Network Redundancy.

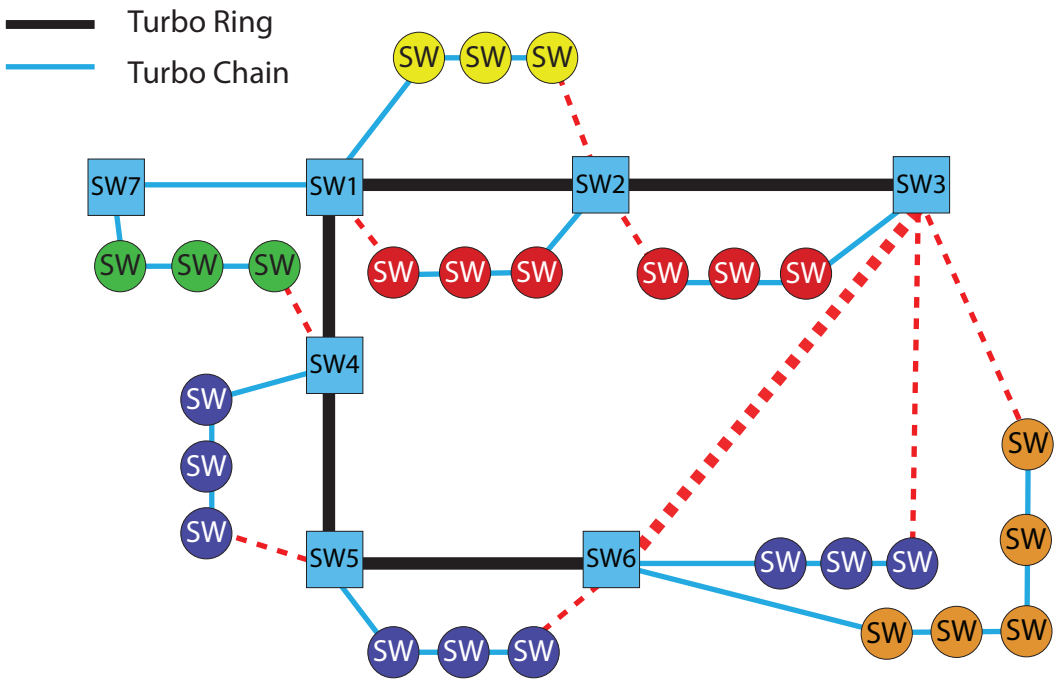

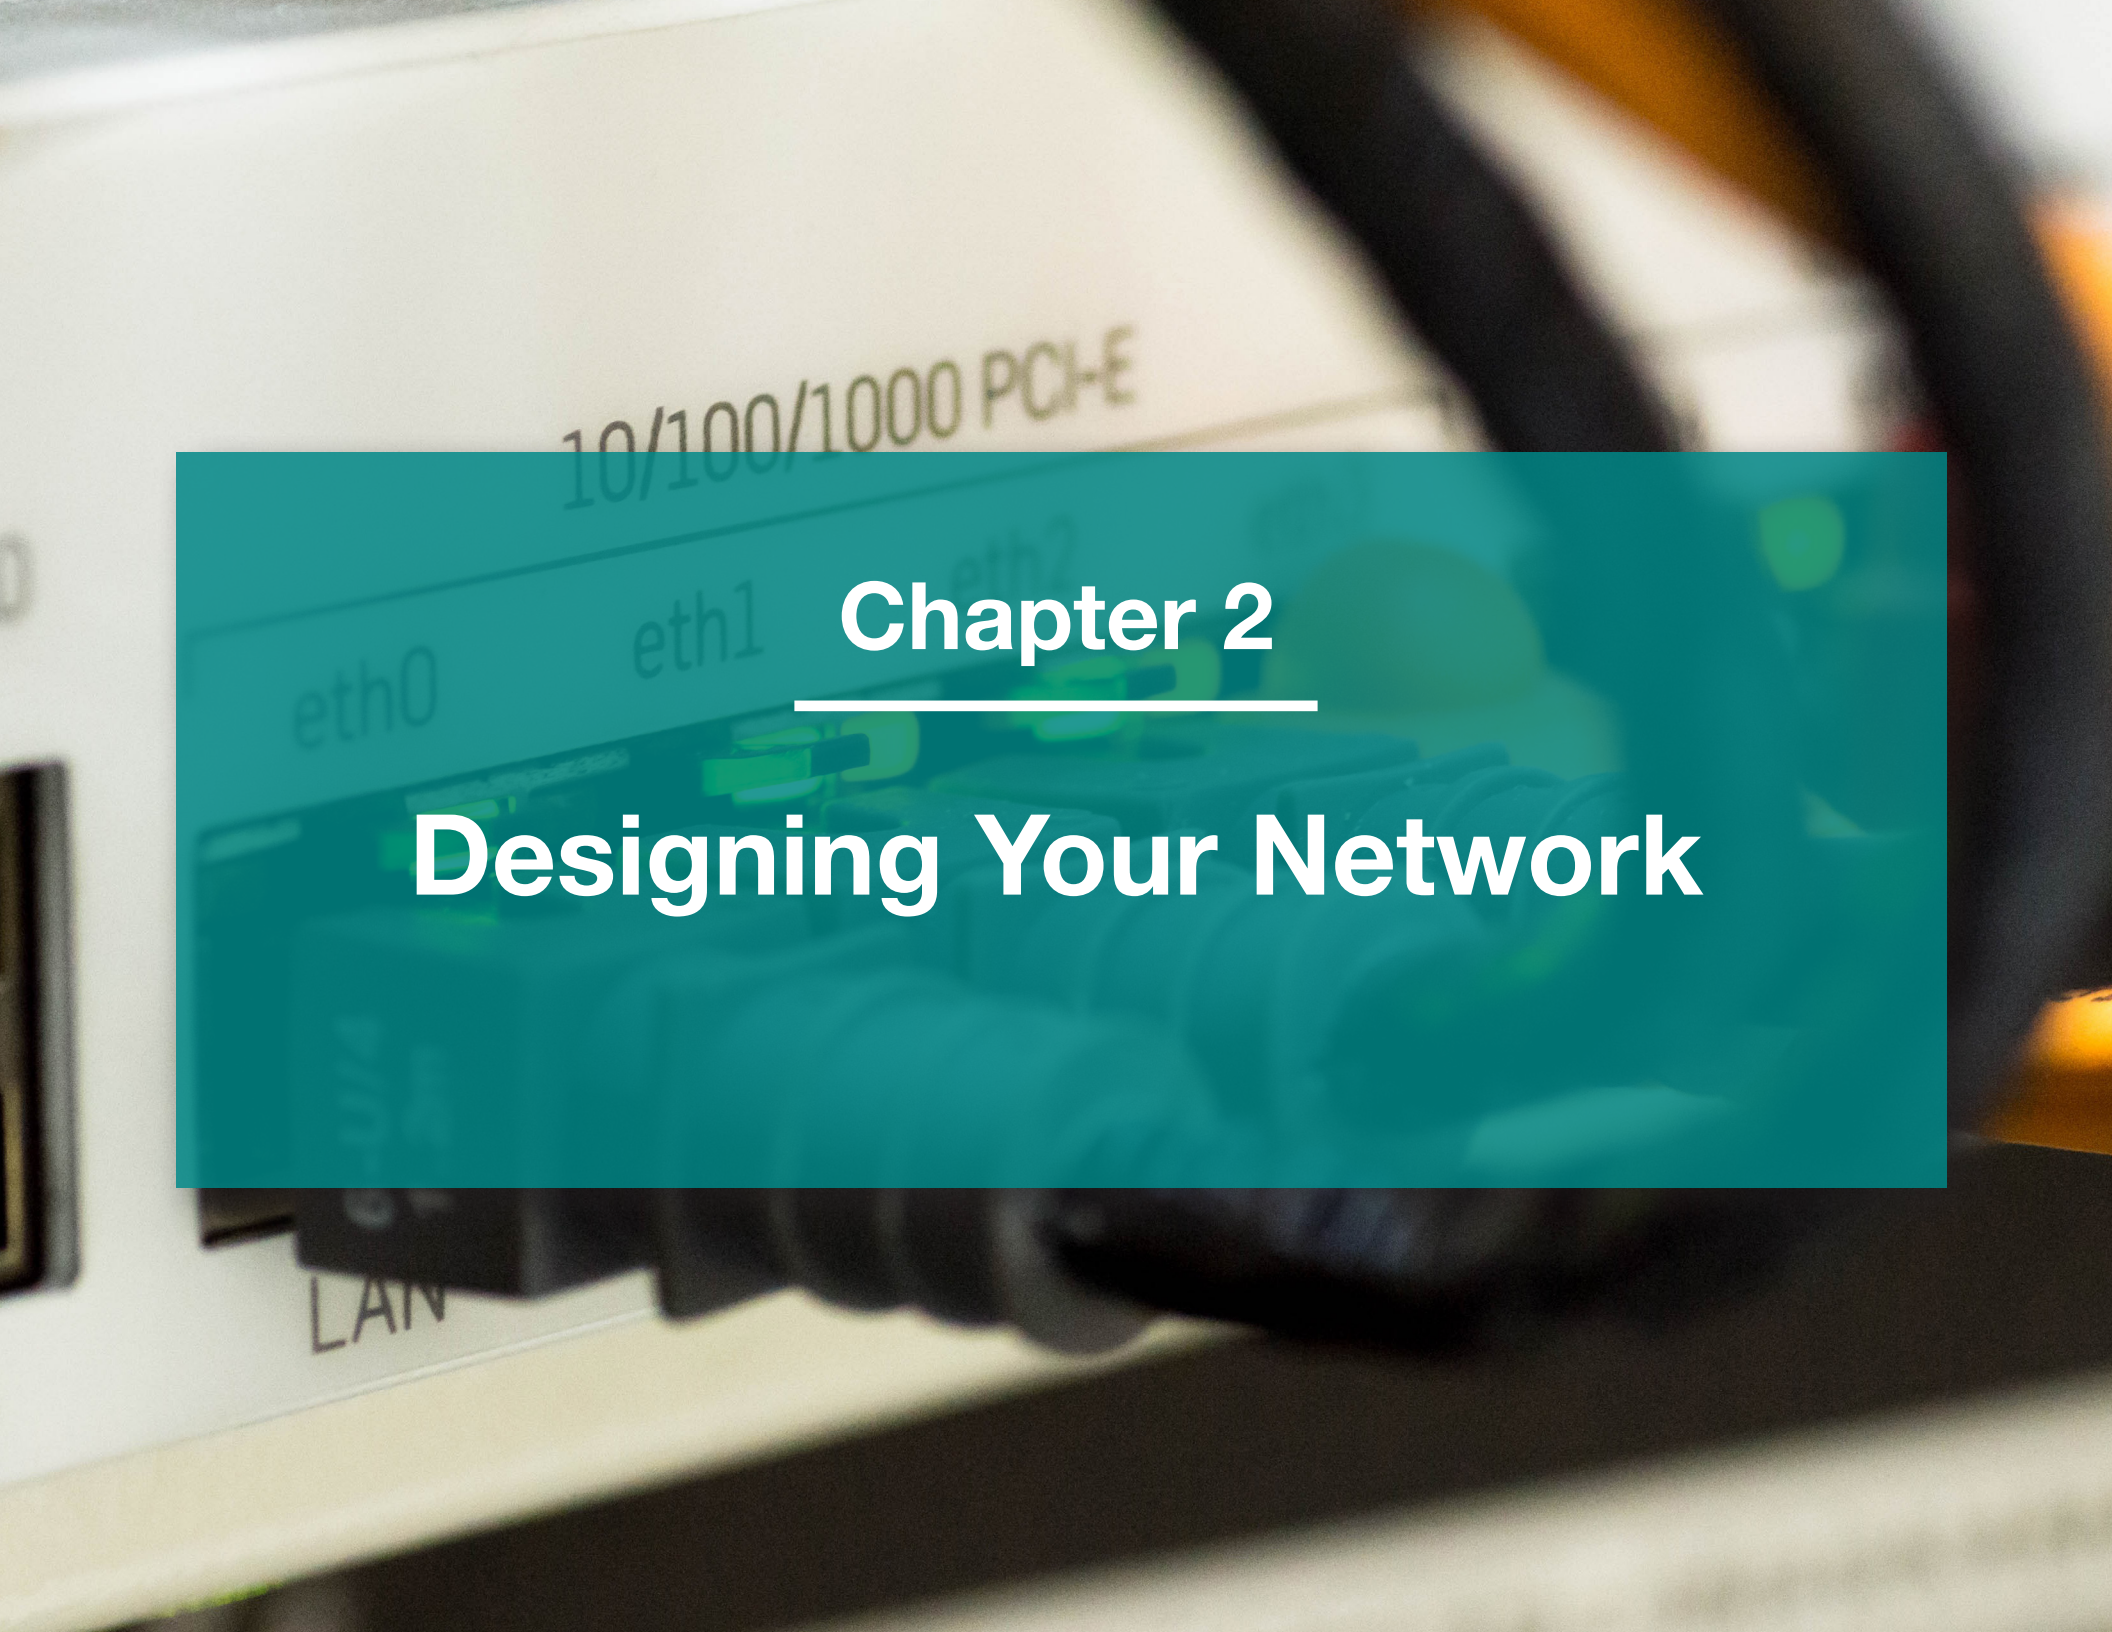

Chapter 2

Designing Your Network

Designing Your Network

Managed or Unmanaged

This question gets asked all too often. The simple answer to the question whether to use managed or unmanaged switches is managed. To be straightforward the benefits to managed switches are so great that it justifies the added cost. Managed switches may cost more at time of purchase, but they have a far lower cost of ownership. With a managed switch, network issues can be prevented before they occur, saving the enterprise from issues that could potentially affect operations. The Return on Investment (ROI) for such a solution is quite high, and can happen in a relatively short time frame. Knowing the cost of downtime, and total cost of ownership can really facilitate making a design choice here.

Download this Payback Calculator
to estimate the payback period for your networking investments by inputting your production and investment costs.

Managed switches allow for network monitoring and reporting, as well as added features to preserve the network in case of failure, such as redundancy, rate limiting, and storm filtering. In many ways, having an unmanaged switch is like driving a car without a dash board - you just don't know what's happening under the hood. When a network issue happens and unmanaged switches are installed, it is very difficult to isolate the problem. It is not uncommon for it to take over a day to recover a network when unmanaged switches are used, and that is assuming the equipment is already on-hand to make the repairs.

The common repair strategy for an unmanaged network is to swap out equipment. First replace the network cables and the switch, then move to the next cable and so on. Obviously this is not a very effective or productive way of resolving the issue. With a managed switch, network issues can usually be avoided before it happens with the switch's reporting capabilities. An NMS (Network Management Software) can be coupled to the solution to expand the ability of viewing the entire network and issues on the network. In the example used earlier with a manufacturing plant that generates $1M in revenue an hour, investing in using managed switches is an easy financial decision. Think of managed switches as a less expensive insurance policy.

Wireline or Wireless

When to choose Wireline or Wireless

Determining when to use wired and wireless can be a challenging decision. Each has a set of advantages and disadvantages. The chart below outlines some of the pros and cons:

	Pros	Cons
Wired	High reliability	High installation costs (when pulling cable)
	High security	
	Unit cost	Many options can be confusing (port count and/or fiber)
	Established/Proven technology	
	Ease of integration with existing networks	Lack of mobility/flexibility
	Availability of products for & to use with	
Wireless	Ease of use in mobile applications	Unit cost higher than wired
	Flexibility/Ease of expansion/relocation	Larger technology hurdle
	Provides long distance/remote coverage	Can be difficult to configure
	Easy/Fast installations	Can have reliability issues if not configured properly
	Can be installed in remote areas where wired cannot	

Once the network requirements are determined, the information can be used to evaluate whether a wired or wireless solution would be a better fit. In general, when mobility is paramount, wireless is the best solution (perhaps the only solution). Cost of installation is another key advantage where wiring and trenching can be a costly endeavor. For wireline, reliability and security is supreme. If the application is mission-critical, wireline is without a doubt the best option. From a security standpoint, it is important to know that unsecured wireless networks can easily lead to data breaches. With wired, you have to be on the physical network in order to even potentially see the data that is being transmitted. If security is of concern, you should use wireline with a secure router or use the most stringent wireless security, called WPA2-enterprise which uses a RADIUS server and certificate-based authentication.

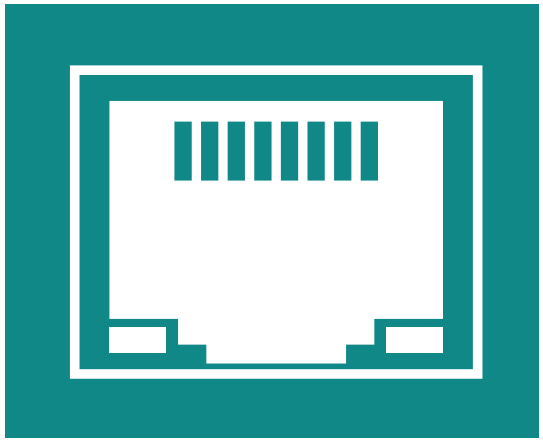

Another important consideration is bandwidth. Wireline offers far better bandwidth than even the latest wireless standards. Keep in mind that wireline is far more efficient and you can get full-duplex operation (the ability to send and receive at the same time). With wireless, efficiency is around 50% of stated throughput and it is half-duplex (can send or receive, one at a time). Now add the fact that the whole wireless channel is one collision zone, whereas with wired communications, the collisions zones can be made obsolete, and system performance can be multitudes of times higher for wireline versus wireless. With that said, wireless does in fact score big when it comes to convenience and cost savings at the deployment stage.

Wireless Best Practices

If wireless is the chosen medium, be sure to do it right and get a site survey. Because wireless uses the airwaves for communication, there are a lot of different ways the signal can be impacted by interference. Site surveys look for interference sources and determines the optimal position for Access Points (APs).

For APs that are adjacent, use channel alternation so that the adjacent APs are using separate, non-overlapping channels. When using the 2.4 GHz frequency band, it is a general design principle to use channels 1, 6, and 11 as these are non-overlapping channels. A non-overlapping channel is a channel in which it does not share bandwidth with other channels. For example, if you were to select channels 2, and 3, each one shares some bandwidth from each other as well as other channels.

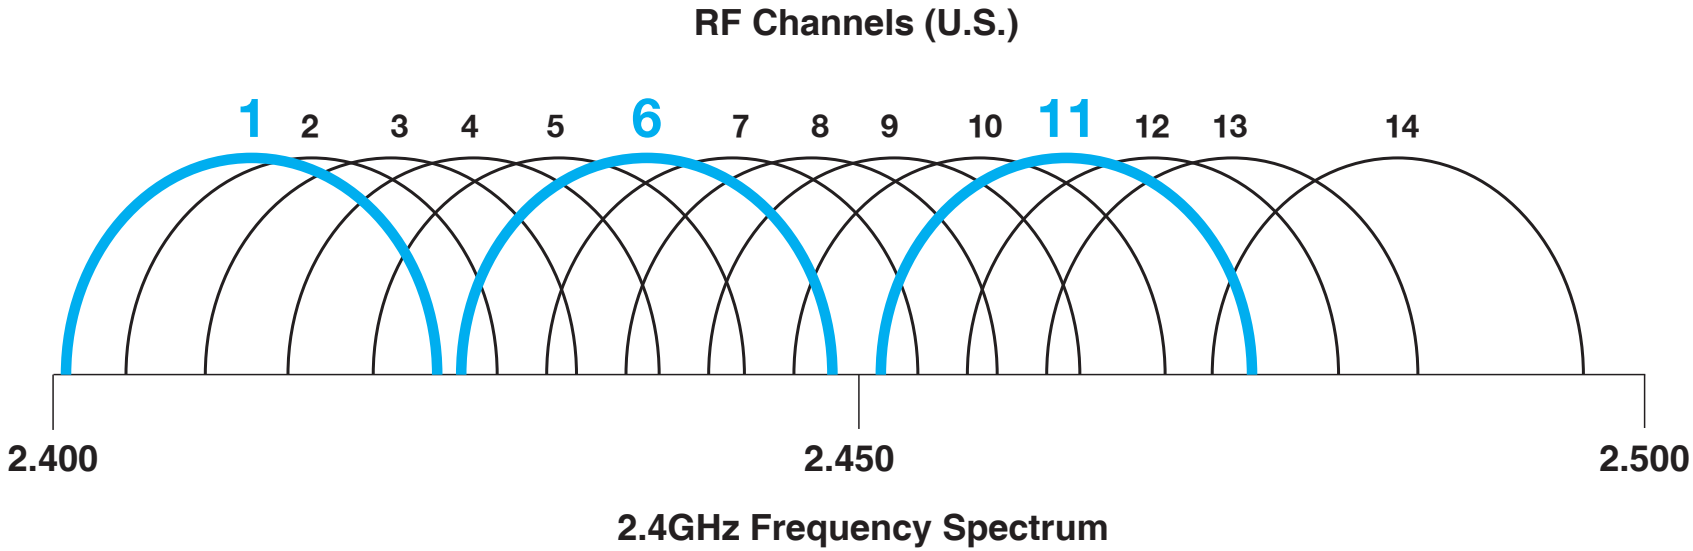

This issue however, is not applicable to the 5.0 GHz frequency band, as each channel is non-overlapping. In addition, it is best practice to have adjacent channels with 20% overlapping coverage to ensure dead spots are avoided for effective roaming capability.

For mobile applications, quick roaming becomes critical. Typically, it can take several seconds for a client to transition between APs. However, there is technology that can cut this process down to the millisecond level either with or without additional hardware like an Access Controller.

Keep in mind that when using APs for a wireless infrastructure, those APs are eventually connected to a wireline-based communication system in most situations. Therefore, wireless and wireline are typically used mutually, with wireless serving as the edge of the network and wireline making up the core and distribution portion of the network.

Wireline Copper vs. Fiber

Even within wireline, there are choices. If wireline is your preference, the determination of when to use copper and fiber will also be a decision to make.

Copper	Fiber
Lower cost	No electrical interference
Supported by most devices	Larger bandwidth capacity
Easy to install	Longer distances
Reliable	Lighter and thinner
Shielded cables are less susceptible to noise	Non-flammable

Copper is much lower in cost, more widely supported, and a lot easier to work with than fiber. These factors make copper a very desirable option for many applications. However, when EMI (Electromagnetic Interference) is a concern, copper will not have the protective requirements to ensure reliability. EMI can be generated by any electronic device from florescent lamps to microwave ovens. In general, the higher powered the device, the more EMI it will output. For this reason, fiber is the preferred choice in mission-critical environments where EMI is high. If fiber is not an option, an alternative is to use copper cables that are STP (Shielded Twisted Pair) in place of the standard UTP (Unshielded Twisted Pair).

Another big advantage that fiber has over copper is distance. Fiber can transmit to much further distances than copper. The design expectation for Ethernet is that you should achieve 100 meters using copper. With STP cabling, you can achieve around 120 meters. With fiber, 20 km or further becomes possible. Since fiber is more versatile than its copper equivalent, it also becomes an area to consider when deciding how to future-proof your network.

	Copper	Fiber
Distance	**100 m** @1,000 Mbps	**40,000 m (40 km)** @1,000 Mbps
Bandwidth	**10/100/1000 Mbps**	**100 Mbps/** **1/10 Gbps**

Network Redundancy

Network redundancy is very important for mission-critical environments. As previously mentioned, the cost of downtime can incorporate several parameters. Therefore, the overall costs of downtime will greatly affect the level of redundancy designed into your network. The following section will cover different forms of network redundancy that should be considered when designing a network. How much redundancy should be built-in depends on the cost of downtime and the mission-critical nature of the application. Bottom line, redundancy is important. Bottom line, redundancy is important [Pun intended].

Power Redundancy

Before jumping into power redundancy, you should understand the basic definition of MTBF (Mean Time Between Failure). MTBF is a metric for the elapsed time before a device is expected to fail. MTBF is usually provided as a specification for industrial equipment so operators have an idea of product life expectancy.

Power redundancy is one of the most fundamental forms of redundancy. Power supplies can fail due to EMI, power surges, or a multitude of other factors. The MTBF of a power supply is generally much shorter than the industrial networking equipment it's connected to. Therefore, it is typically expected that the power supply will be the first physical device to fail. Most industrial equipment have the ability to support more than one power supply which enables power redundancy in case of a failure. When tied to a managed switch, the switch can be configured to alert the user of a failure. This alert can be network based, such as an email, a notification appearing on the NMS, or it can be as simple as a flashing light using the included DO (Digital Output) port. Fortunately, the loss of a power supply does not need to affect network communication in any way and is easy to implement.

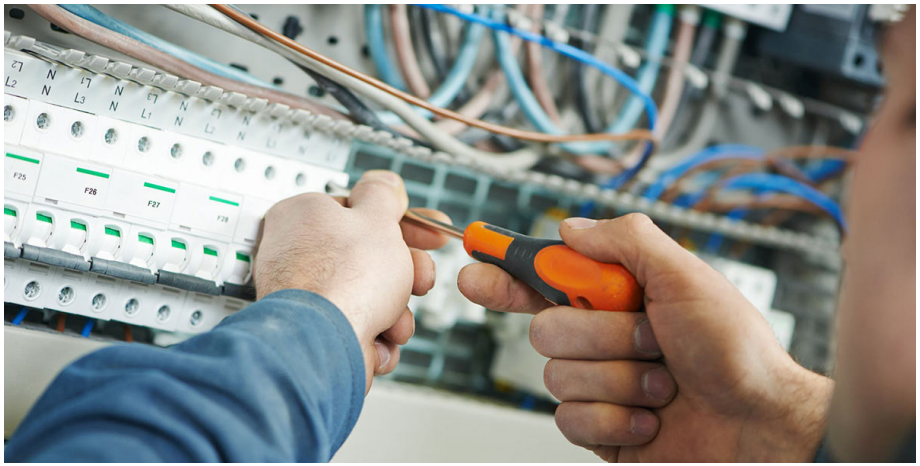

Safeguarding against power failure is also important and can be achieved through the use of a UPS (Uninterrupted Power Supply). There are many different types of UPS's to choose from. Generally, a battery-less system that uses a capacitor is the most reliable setup in the industrial setting. However, this method will not filter out noise or spikes on the power line. A battery version will usually condition the power to ensure it is clean, resulting in longer life of the power supply and connected equipment. Unfortunately, this is usually not as hardened as the capacitor equivalent. In either case, a UPS will safe guard in case of power failure.

Layer 2 Communication Redundancy

The truth is that Ethernet was not originally designed with redundancy in mind. If you have an unmanaged network, do not try to connect your switches in a loop or ring topology. The result would be a switch loop which will in turn generate a broadcast storm and MAC address flooding that will quickly bring down your network. In fact, we have seen this happen at a location where a technician saw the end of an Ethernet cable lying around and thought it was accidentally disconnected. The technician connected the loose cable to the unmanaged switch, resulting in the organization of around 1,000 employees to lose its network for over a day and a half. The issue was not able to be resolved until the link that created the broadcast storm was traced down. This type of situation can easily be avoided by using managed switches that support redundancy protocols.

Creating a fail-safe network is essential in a deterministic network. What if a cable breaks, gets disconnected, cut, or if a switch loses power, has an internal failure, gets water damage, or if the power supply fails? There are countless ways for a failure to occur. Having a backup creates more network resilience; however as stated earlier, if this is attempted with an unmanaged switch network the result will be an immediate failure of the "entire" switch network. Ethernet does not allow for loops, and so this is where redundant Layer 2 protocols like RSTP (Rapid Spanning Tree Protocol), Turbo Ring, and Turbo Chain are needed.

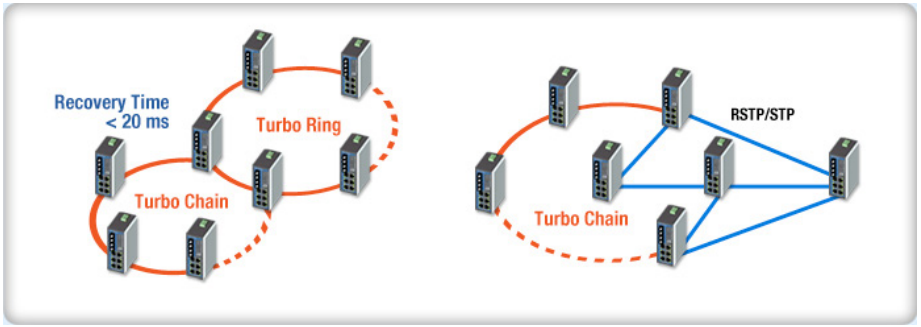

A redundancy protocol allows Layer 2 networks to be connected into rings, and in some cases, full meshes. It does this by logically disabling a port so that no switch loops are formed. If an issue like a link drop or a switch failure occurs, the blocked path begins to forward traffic – this is when the network recovers itself.

How fast the network needs to recover is based on the environment it will be installed in. The most common protocols in use today for Layer 2 redundancy are STP (Spanning Tree Protocol) and RSTP. These two technologies have the advantage of being a standard and for allowing a wide array of topologies. Unfortunately, these technologies were designed to be used in the commercial sector where 30 seconds of lost email communication is not a big concern. We have already explored the differences and mission-critical nature of the industrial sector. Just keep in mind that in the example of a $1M per hour manufacturing plant, the cost of downtime is $0.28 per millisecond.

There are two ways to reach recovery times in the milliseconds, in a reasonably sized network. The first is to use proprietary technology, such as Moxa's Turbo Ring and Turbo Chain. The second is to use cutting edge technology like PRP and HSR. The benefit of the Turbo Ring and Turbo Chain technology is that it offers a recovery time of 20 ms in a network of 250 switches, and the technology exists in each Moxa managed switch. Another advantage is that multiple Turbo Rings and Chains can be used in a single Layer 2 network, and each Ring and Chain can have independent convergence zones. To explain what this means, it is first important to explain what convergence is. Convergence is the act of recovery in a network when an issue is found. The typical convergence time for STP is 30 seconds, RSTP is less than 1-5 seconds, and Turbo Ring/Chain is 20 ms.

With STP and RSTP, the entire Layer 2 network forms what is known as a tree. A network issue like a downed port means that the entire tree will converge. For a flat network, that means the entire network will converge. When the network is converging, all communication on the network stops. In the case with STP, it means that network communication will stop for around 30 seconds for each incidence. With Turbo Ring and Turbo Chain, only the individual ring or chain segment will converge, not the entire network.

In the topology below, the black and blue lines are the primary paths, and the dotted red lines are the backup paths. The ring (black lines) and the chains (blue lines) each form independent convergence zones, each with a recovery time of 20 ms or less.

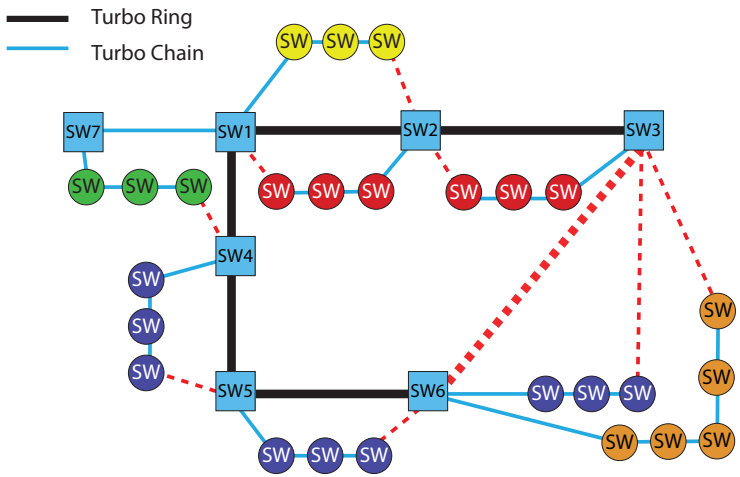

PRP and HSR is another technology that has the benefit of being a standard. Its main benefit is that it offers 0 ms recovery times. It accomplishes this task by creating redundant traffic. The disadvantage of this technology is that it will cost a lot more in terms of network infrastructure and it is dependent on purpose built devices known as RedBoxes (you cannot rent a DVD from this one). A RedBox is also known as a redundancy box which offers both PRP and HSR redundancy. PRP can easily double the cost of the network, where HSR will limit the network to only work with purpose built devices and limit the amount of convergence on the network.

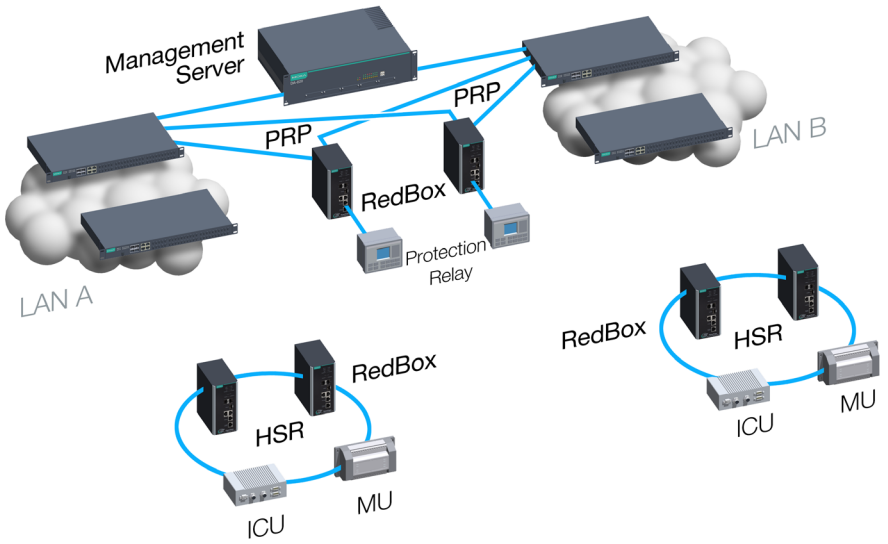

Network Segmentation

Networks can easily become too large, especially Layer 2 networks. A simplified explanation of a Layer 2 network is a network which can communicate without routing. There is more to it, but for the purpose of understanding, this definition works best. Network segmentation is breaking up a network into Layer 2 partitions. There are multiple ways to achieve this. The most common is via VLAN (Virtual Local Area Network) and Layer 3 switches or routers.

The benefits to network segmentation are numerous. One of the primary reasons one will do this is that it will improve performance in a large network. The reason for this is that it breaks up the broadcast domain. To best explain what the broadcast domain is, draw this image in your head, a room with people in it. In the room, you have people all around talking. When one person broadcasts, it is analogous to a person shouting to the group. You can even picture that person having a bull horn. Everyone in that room has to hear you (listening to you is another story). They may not be interested in what you have to say, but they are forced to hear you. It would be impolite to talk over you. Now picture that same scenario in a much larger room. Each time someone chooses to broadcast out, everyone is forced to hear that person, and stop their communications.

Networking works similar to this example. Broadcasts are actually quite common as there are many critical network processes that rely on it, such as ARP (Address Resolution Protocol). When the broadcast domain is small as the one pictured, communication still works well. With 15 devices, and each device broadcasting out, there will be a total of 14 broadcasts.

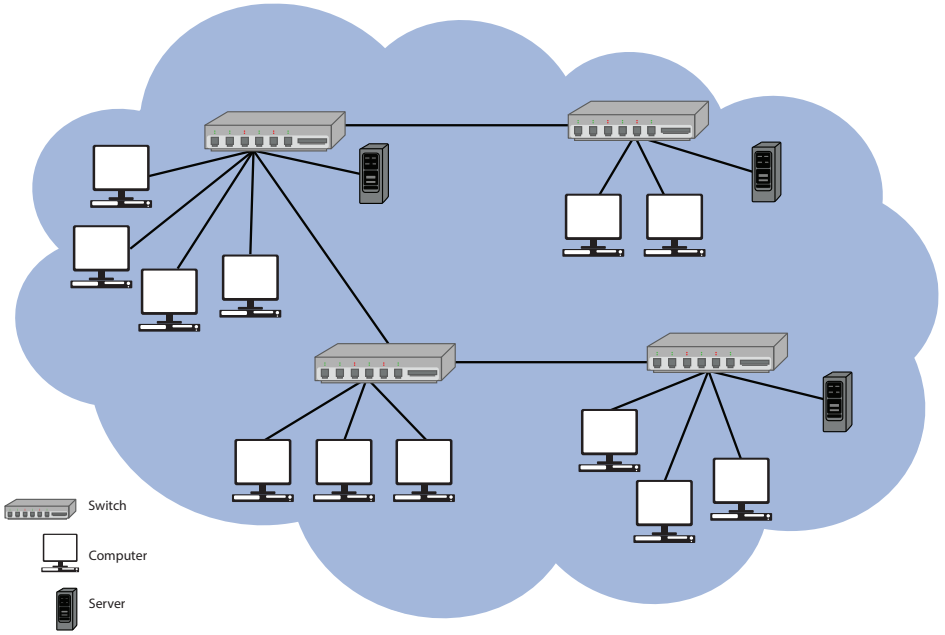

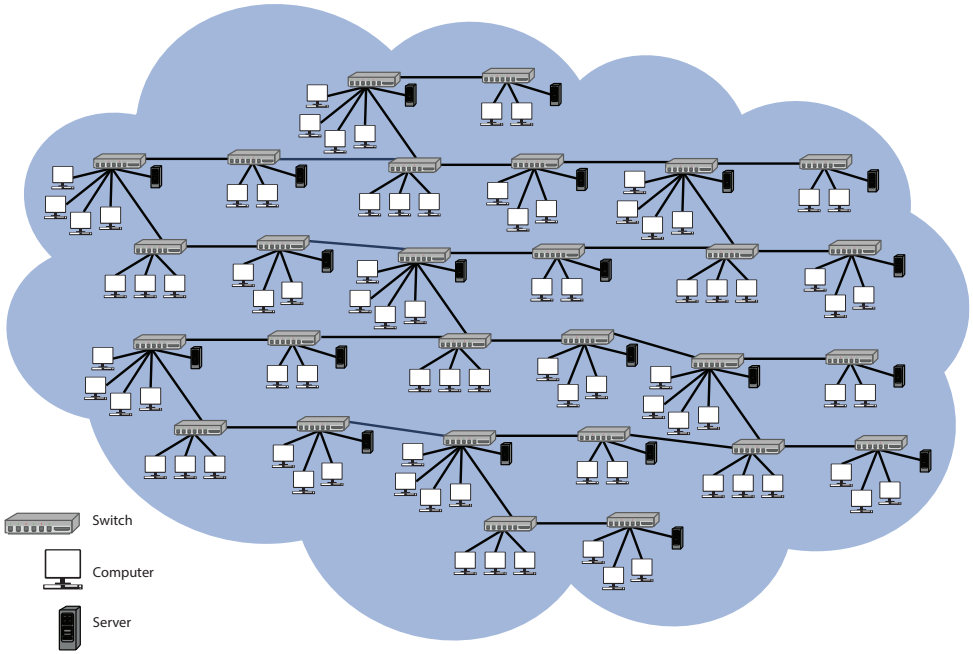

Now increase the scale to a much larger network with 1,000 devices. If everyone sends a broadcast destined to 100 devices, the network becomes much busier. This is not as extreme of an example as you may think. Broadcasts in networks are common place, and so limiting the broadcast domain increases network performance and increases system stability and reliability.

The other advantage of network segmentation is increased security by confining and restricting network traffic. For example, traffic on the automation floor can be restricted to only the plant floor and select key decision makers. In addition to keeping selected parties within the network, segmentation also helps to keep unwelcome visitors out. For example, if a hacker were ever to successfully access a network partition, he will not have access to critical systems depending how policies are set. Network segmentation is similarly used at the enterprise level when traffic is restricted from an HR network, for example. However, common systems like email servers and such may have less restrictive policies in place.

VLAN

Networks can be segmented by function or location. Location was the most common form of segmenting until recently. With the advent of Virtual Local Area Networks (VLAN), function based network segmentation has become the preferred method of network segmentation.

A VLAN, is a Layer 2 network that is logically segmented. What this means is that a traffic slowdown from a given network is segmented and kept separate, which would also be the case if separate networks were created.

The key advantage of accomplishing this task with VLANs as opposed to dedicated networks is cost savings and ease of management. In the past, a whole network with its own dedicated switches and infrastructure would have to be created for each traffic flow you needed to isolate. With VLANs, the same degree of separation can exist on a shared network infrastructure.

There are two primary types of VLANs: Tagged and Port. Port based VLANs are only useful in small networks, and have the advantage in which traffic is not modified in any form. Tagged based VLANs following the industrial standard 802.1Q, can work in networks of any size as well as cross vendors. 802.1Q tagged based VLANs have become the most predominant form of VLANs due to its scalability. It is the recommended type of VLAN for most to utilize, unless there are only one or two switches in the network. In this form of VLAN, Ethernet frames are modified. Most end devices will not be able to read these modified frames, however, managed switches have the ability to add (push) and remove (pop) VLAN tags. In this case, the end device will be unaware that it is on a VLAN. These devices unaware of the VLAN will still have all the benefits allotted by VLANs with the managed switch handling the VLAN interaction in the backend.

This figure below illustrates two tagged based VLANs. VLAN1 is shown in red and VLAN2 in blue. The end devices may not be aware that they are on a VLAN as the switch will add and remove the tags as soon as the Ethernet frames reach it. For example, in the case of the computer marked “A”, the Ethernet traffic coming from this computer will be untagged, or not containing any VLAN information. When it reaches the switch, the switch will “tag” the frame under VLAN 1 (the blue VLAN). The network can be designed to pass the VLAN tags from switch to switch. Even though the same network infrastructure is used, there are two separate network segments. Without a Layer 3 device, these network segments will not be able to cross communicate.

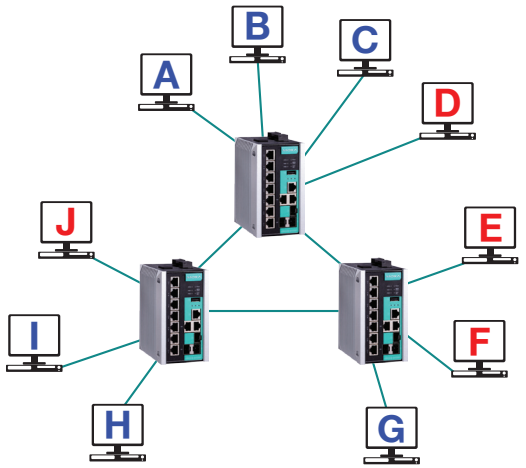

To put it more simply, if VLAN 1 and VLAN 2 were created, then a device in VLAN 1 can communicate to any other device on VLAN 1. Likewise, a device on VLAN 2 can communicate to any other device on VLAN2. However, devices on VLAN 1 cannot communicate with devices on VLAN 2 directly. This includes broadcasts, thus using VLANs is an effective way of breaking up the broadcast domain.

Now that we've established that devices on VLAN 1 cannot communicate with devices on VLAN 2 directly, what if the need arises where communication between VLAN 1 and 2 is required? Usually, this will be done with the introduction of a Layer 3 device. When introducing Layer 3, it is best to use a firewall or ACL (Access Control List) to limit the amount of cross-communication between the different network segments.

Another requirement for properly designing a VLAN network is to ensure that each VLAN has its own network address (commonly but mistakenly called its own subnet). It is also a common mistake to use the same IP address for an entire VLAN network. When the network is segmented, there won't be any issues as switches use MAC address, not IP address to make its forwarding decisions. However, the problem is when Layer 3 networks come in play. In order to route between the different VLAN segments, Layer 3 networks need to determine the boundaries of these networks. This is done by ensuring each VLAN has its own network address.

Layer 3 Switches and Routers

A Layer 3 device will be used when different networks need to be bridged. These different networks can be physical networks as well as VLANs. Within a network VLANs are usually the preferred method of network segmentation. When a network needs to integrate with the internet, also known as a WAN (Wide Area Network) based network, then VLANs will not work as the WAN is physically different than the LAN (Local Area Network). In other words, Layer 3 is the opposite of network segmentation. Network segmentation separates traffic flows, whereas Layer 3 connects the disjointed traffic flows together. Even more so, it joins the traffic in a controlled manner. A Layer 3 device, whether it is a Layer 3 switch or a router serves one primary purpose, which is to route traffic.

The table below shows the key differences between Layer 3 switches and routers. One thing to keep in mind is that a Layer 3 switch for all intents and purposes is a router. More specifically it is a hardware based router, where traditional routers usually accomplish its routing via software. Hardware routing means that it can route much faster than a software based router. This is because hardware can do simultaneous tasks in real-time whereas most software is sequential. There are trade-offs. Executing tasks with hardware is more difficult, therefore Layer 3 switches will not have all functions implemented in hardware. Firewalls for example, is a software feature that is usually absent on a Layer 3 switch. It is possible to put a firewall in a Layer 3 switch, or any other software based features, but it will create a bottleneck that will hamper overall system performance.

	Layer 3 Switch	Router
Routing Performance	Extremely Fast	Slow
Port Density	High	Low
Network Routing Capability	VLAN Based	Interface Based (Usually)
Virtual Private Network (VPN)	Not Capable	Fully Capable
Network Address Translation (NAT)	Not Capable	Fully Capable
Firewall	Access Control Lists (ACL)	Fully Capable
DHCP Server	Limited (Usually)	Fully Capable

Layer 3 switches do have ACLs in which traffic filtering can be done at the hardware level. It is not as robust as a router, but basic filtering and access control can be accomplished. For this reason, routers are almost always found between the LAN and the WAN, such as the Internet. Routers have far more capabilities than a Layer 3 switch, however for networks that do not connect to the Internet, a Layer 3 switch with its high-performance routing is ideally suited.

Edge to Core

The size and scale of a network determines the overall topology. When networks expand, it is a best practice to break it up into a hierarchical design. The three hierarchies are edge, distribution, and core where each layer serves a unique purpose. Not all three layers will be required in every network. A small network may be able to get away with only an edge layer network. As it scales, distribution is introduced. If it becomes even larger, the core layer is added.

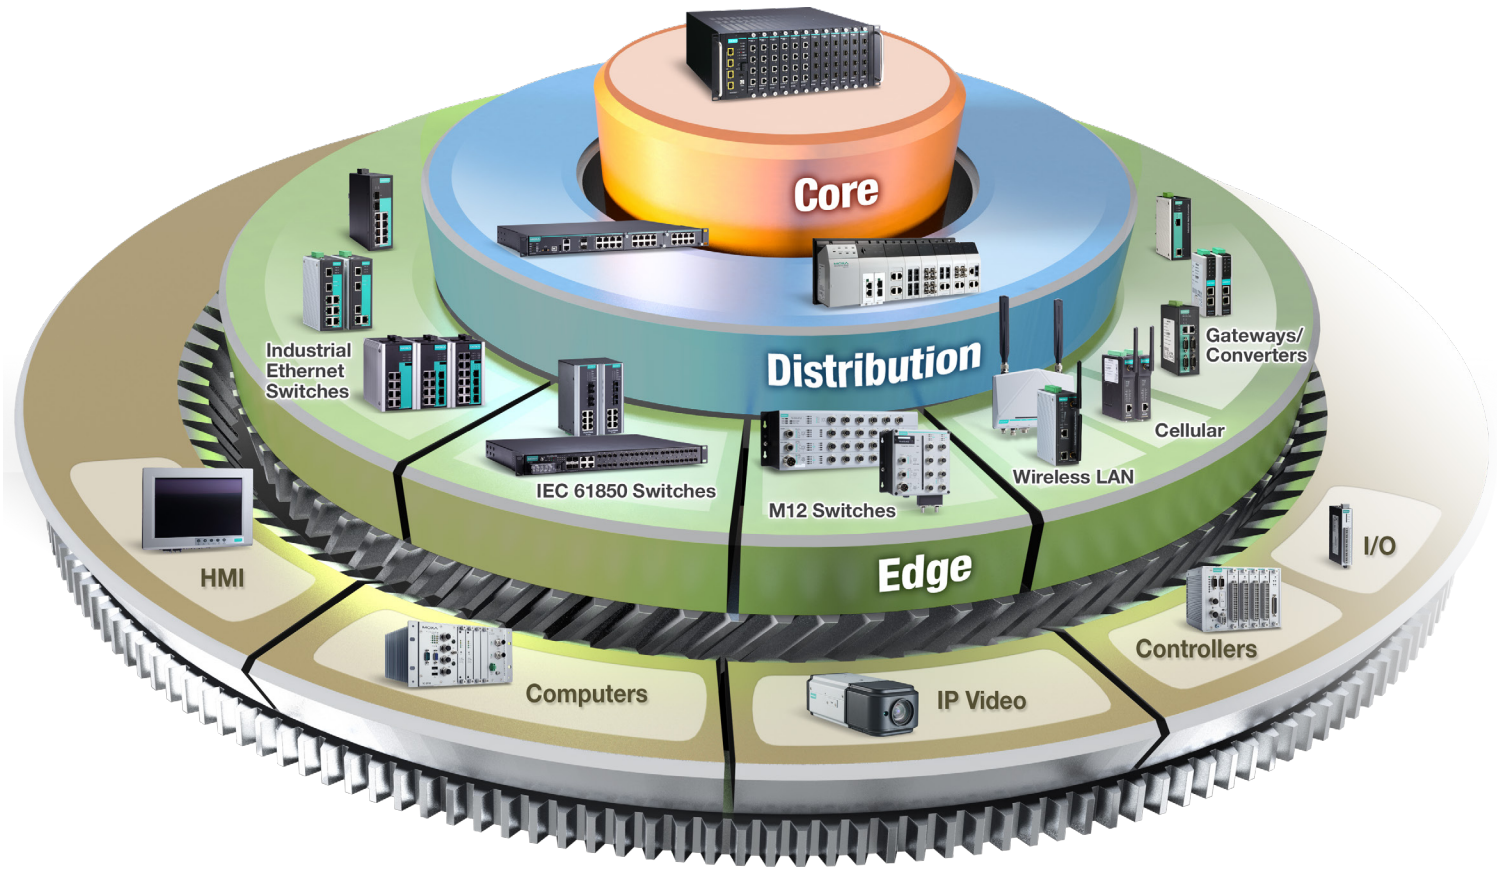

Edge

The edge is the point in which the network touches the end devices, whether those end devices are PLCs or PCs. Therefore, the edge layer is the layer closest to the end devices and the one users have the most interaction with. Because these devices are at the edge, there is little need for high port counts. In the commercial/enterprise space, it is common to create a star topology in which all links go to a central point. This does not provide any redundancy and if the aggregation switch fails, there will be a loss of many end devices. Since uptime is of importance for industrial networks, they are typically distributed in nature, and switches of lower port count are deployed at the edge. These edge switches can employ Layer 2 redundancy technologies to prevent loss of traffic from any individual link loss. Also, its distributed nature ensures there is not one single point of failure in regards to the network infrastructure.

Distribution

The primary role of the distribution layer is to join all the edge layer switches into a larger network fabric. This is where network segmentation and Layer 3 routing usually takes place and where Layer 3 switches are often the preferred device because of its fast routing and high-port density. The combination of edge and distribution serve most mission-critical networks.

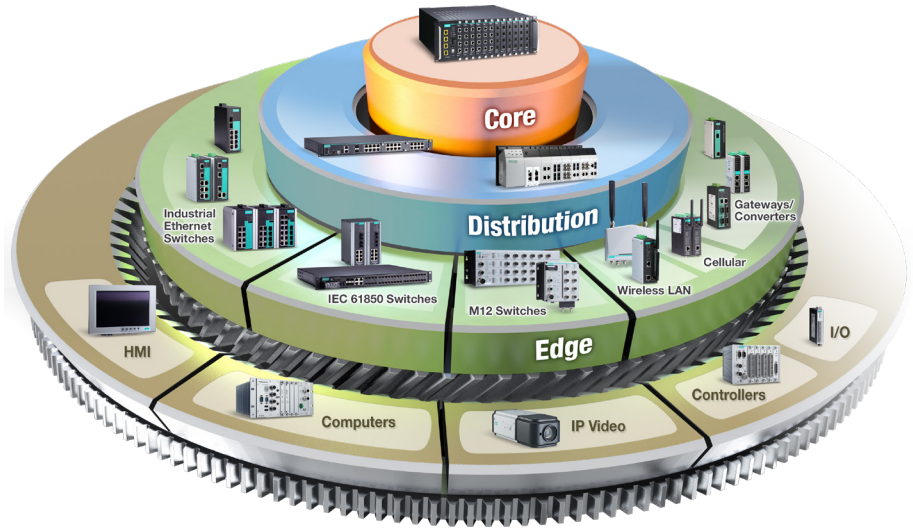

Core

Core switches are the most central part of a network. These switches have high speed links such as 10-Gigabit per second and are able to route quickly using Layer 3 switches. These are usually deployed in large networks that have demanding requirements such as video. In a network that uses all three hierarchies, core switches are usually Layer 2, as the routing is handled at the distribution layer.

To Multicast or Not To Multicast

Before taking a deep dive into best practices of multicasting, let's start with a brief explanation of the underlying technology. There are several different forms of Ethernet communication: unicast, multicast, and broadcast. Essentially there are one-to-one, one-to-several, and one-to-all forms of communication, respectably.

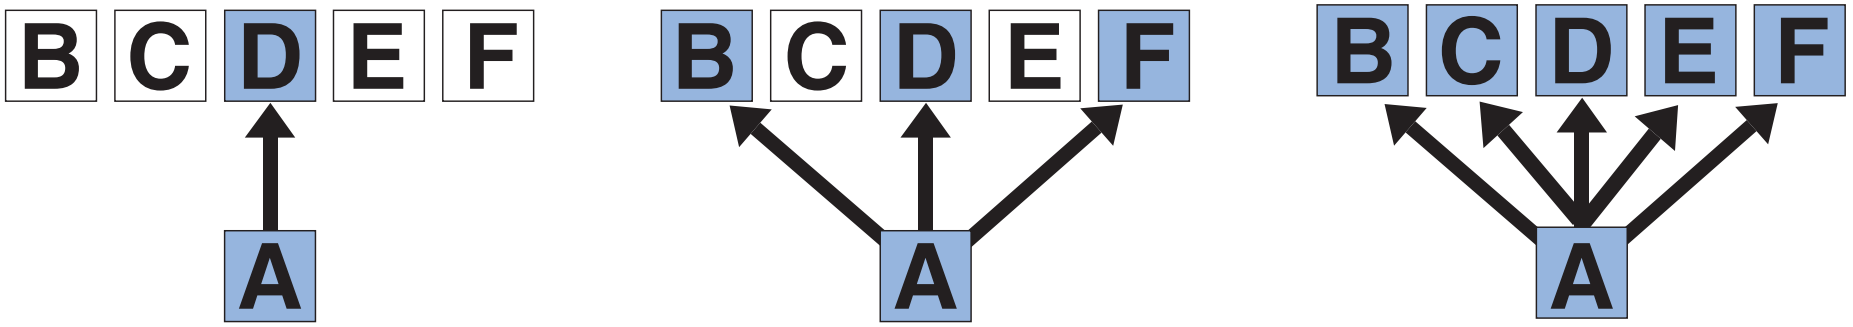

Unicast:
One to One

Multicast:
One to Several

Broadcast:
One to All

The chart above provides a better explanation on the three different forms of communication. With unicast (one-to-one), only one message is going to one receiver. With broadcast, all receivers in the given broadcast domain are receiving the message. The challenge is when the need arises to select several devices that need to receive a message. This is where multicast (one-to-several) is used. Multicast is more difficult to accomplish when compared to unicast or broadcast.

Multicasting is group communication. There are multiple methods to accomplish multicasting, but the most prevalent one is to use multicasting protocols like IGMP (Internet Group Management). IGMP itself is a Layer 3 protocol designed for multicast routing. Switches are Layer 2, thus do not support IGMP natively. IGMP helps manage the groups as well as forward the multicast data streams efficiently.

Rather than support the full protocol, switches implement a multicasting method that is known as IGMP Snooping. The switches “snoop” or eavesdrop on the IGMP communication process. The process in which IGMP Snooping works is not standardized, and every manufacture could have a different method to accomplish the same goal. However, IGMP Snooping requires switches to have some form of intelligence, thus a managed switch is required. Unmanaged switches will process IGMP multicast frames the same as all unknown frames, which is to flood. Flood essentially means that the switch will broadcast out the message, which means it becomes one-to-all. Multicast transmission affords many advantages over unicast transmission in a one-to-many or many-to-many environment:

Greater Efficiency: available network bandwidth is utilized more efficiently since multiple streams of data are replaced with a single data flow. As a result, you achieve greater bandwidth utilization.

Greater Performance: fewer copies of data that require forwarding and processing.

Distributed Applications: multipoint applications will not be possible as demand and usage grows because unicast transmission will not scale (traffic level and clients increase at a 1:1 rate with unicast transmission).

Scenarios in which multicasting is heavily used is video, audio, and Ethernet/IP installations. The first question to ask before deciding whether to use IGMP is: **Does the end device support it?** Most VoIP (Voice over IP) and video cameras/encoders support IGMP and it is part of the standard specification for Ethernet/IP devices.

The second question is: **Is there a need?** For example, just because there is an application utilizing video and the video end device supports IGMP, does not mean that it is the preferred method of communication. There has to be multiple receivers (or consumers) for the transmitter (or producer). If there is only one receiver with several transmitters, there is no benefit to using multicasting. However, if there are several receivers per transmitter, there is definitely a need. With Ethernet/IP devices, the rules are a little different as each device can be a producer and consumer of multicast. If there is more than one device on the network, and you are using implicit communications, then it is a good design principle to use multicasting.

Ease of Management

A commonly overlooked decision in network design is how to run and maintain the network. Identifying who is going to maintain the network is critical to successful network longevity.

Here are three areas to consider when it comes to network management.

Complexity: Is your network design too complicated? Adding complexity to the network makes it more difficult to understand, especially to those who need to manage it. That said, there are technologies that can help prevent this issue from starting in the first place, thus a delicate balancing act would have to take place.

Talent: Does the staff responsible for network maintenance have the skill set to manage the network at its level of complexity? Keep in mind that these skills can always be acquired through proper training. In these instances, it is important to choose a solution provider that will not only sell the equipment, but also be able to train on and support it effectively.

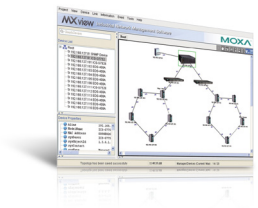

Systems: Did you invest in management software? An NMS (Network Management Software) is a software utility that helps monitor network health, and alerts whenever issues arise. In some cases, issues can be averted entirely. Investing in an NMS makes a big impact on simplifying network management.

An NMS is usually vendor neutral, assuming the NMS supports the importing of MIBs (Managed Information Base), files used to monitor network equipment. Another requirement is that the network gear to be monitored supports SNMP (Simple Network Management Protocol).

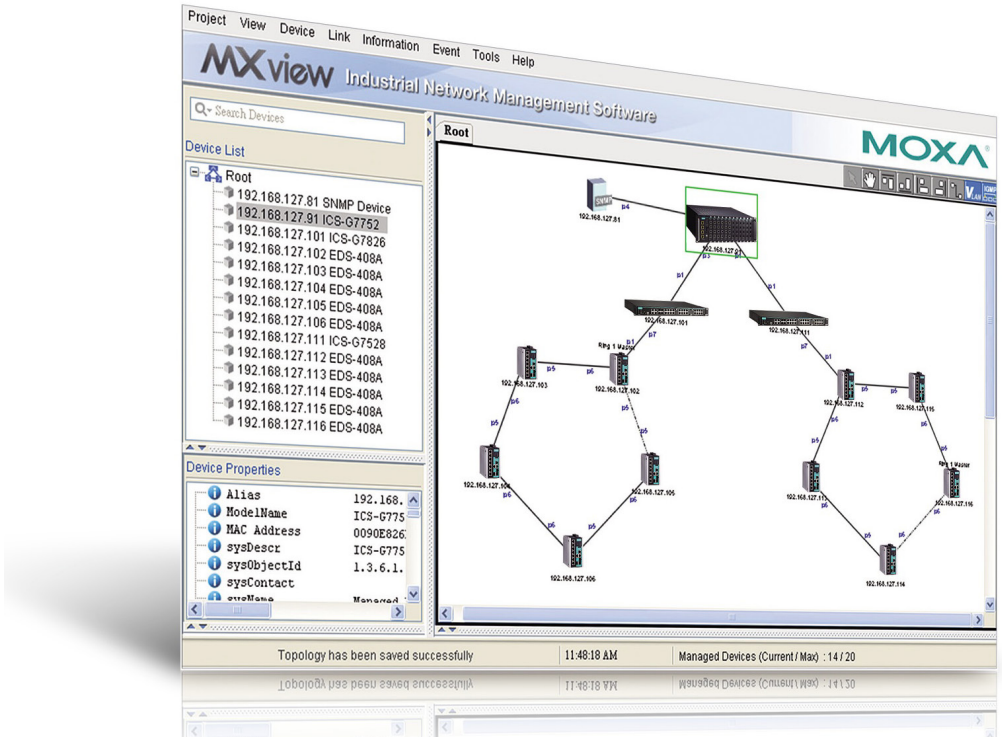

There are some advantages to using vendor specific NMS solutions. For example, with Moxa, the NMS not only monitors network gear, but can configure and troubleshoot it as well. This means that the software suite can be used from the commissioning of a network, to its monitoring, and even help with troubleshooting the network. Currently, there is no solution that is not vendor centric in which all the above can be achieved. Another advantage of Moxa's NMS is that most NMS are designed for the commercial space. Using an NMS for the industrial space allows for a more robust and easy to read monitoring solution designed for mission-critical applications, which are not always going to be monitored by IT professionals.

Download our white paper on:

Overcoming Obstacles of Industrial Network Management

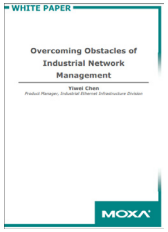

Security

Security is extremely important when designing an industrial network. The issue is that industrial networks are prime targets for hackers and would-be thieves. Hackers can affect important operations and steal intellectual capital. To overlook security can be a dangerous move. In adopting security measures that protect the entire network, the critical considerations include:

1. **Protecting against EXTERNAL hackers who want to steal information**
2. **Protecting against INTERNAL employees who accidentally cause a disruption**
3. **Protecting against INTERNAL hackers who want to cause a disruption**
4. **Protecting against INTERNAL employees who want to steal information**

Firewalls provide protection against hackers on the Internet.

Protecting against EXTERNAL hackers who want to steal information

Recently, it was reported that intellectual property theft costs the United States as much as $300 billion and 1.2 million jobs annually (Commission on the Theft of American Intellectual Property). The majority of this lost revenue can be attributed to a lack of security in our networks and communications systems. In fact, a well-documented report titled APT1 by Mandiant shows that a significant amount of IP was stolen from industrial control systems that were network-enabled. In many cases, cyber-attacks involved hacking into a network and downloading CAD files for very specialized patented designs. Once the files were stolen, the product plans could easily be copied, allowing imitation products to flood the market, destroying the advantage innovative companies gain through large investments in R&D.

In most cases of intellectual property theft, the IT network and the industrial automation network were tied together. It is equally likely that the security policies between the two groups diverged. To hedge your bets, you should at the very least use a firewall/router to separate the IT network from the automation network. No one from the office side should have access to proprietary design files, even IT administrators. For example, would you allow an industrial engineer to access employee personnel files from the HR department?

Your second layer of defense is to make sure any WAN connection also has a firewall separating the WAN from the LAN side. Additionally, remote access should only be allowed by using at the very least, a 128 bit encrypted VPN to protect your data and manage access to your industrial control network. Besides incorporating encryption and authentication, VPN access should also be subject to strong user passwords to prevent unauthorized access to your network. Such passwords can be derived using a password generator that helps create passwords that conform to best standards. This is the primary means of setting passwords, and is recommended by the NIST (National Institute of Standards and Technologies).

Protecting against INTERNAL employees who accidentally cause a disruption

Once you have shored up your defenses to the outside world, you will want to address users within your own company. Unintentional hacking is probably the most common cause of network disruption and is often caused by devices that are not configured properly. Most employees are not extremely tech savvy and the last thing they want to do is harm anyone or any system within their own company. Despite the absence of malicious intent, it is still best practice to use simple subnet segmentation. In other words, each cell should have its own small, protected network. More specifically, this network should prevent accidental entries by requiring purposeful intent to access. Furthermore, segmentation protects against broadcast storms or broadcast data produced by a misconfiguration outside of that cell. Segmentation can be achieved with a simple router in each cell.

Protecting against INTERNAL hackers who want to cause a disruption

Although the vast majority of employees do not wish their employer any harm, there are always exceptions. For example, disgruntled contractors or recently laid off employees may hold a grudge. To prevent any critical device interruption, it is good practice to place industrial firewalls in front of your most important equipment. An example might be to place a firewall between the rest of the network and a manufacturing cell, or a group of critical control PLCs. By placing a firewall at this section of your network, you can ensure that only industrial control messages get through. For example, you can configure the firewall to only allow EtherNet/IP communication, and only from specified senders. Although this would normally be accomplished by simply segmenting the network with subnets, this method alone is not nearly as effective as incorporating both segmentation AND firewall protection.

Protecting against INTERNAL employees who want to steal information

In addition, disgruntled employees may want to take some parting gifts as they leave a company for a competitor. In this case, the stealing of intellectual property may result in the loss of significant revenue. Proprietary design files are very valuable and could bring a company to bankruptcy if they end up in the wrong hands. Here, a firewall in front of your critical devices that is capable of DPI (Deep Packet Inspection) could really save the day. DPI basically looks at the actual data itself (which could be specific commands and requests) to determine whether it should be allowed to pass through the firewall.

Example: A malicious industrial engineer is trying to steal a design file from a CNC/DNC machine. He sends a MODBUS request to download this file. If configured properly, the firewall would see that he is not just requesting machine status as usual, and would deny his request. Allowing basic non-harmful commands may be necessary on a day-to-day basis. But with deep packet inspection, you can protect against unauthorized engineers issuing commands they should not be using. Some industrial firewalls already have DPI capability, which can be best utilized if deployed.

Best Practices for Increasing Network Security

There are countless devices connected to industrial control networks, and these devices also play an important role in your security policy. From a best security practice standpoint, it is important that each connection follows the guidelines listed below:

Disable Telnet on any network attached device (make sure you have another way into the device first). Use SSH instead if available.

Telnet transmits characters in plain text. If someone is using a software network sniffer, this can be a major security risk. SSH connections accomplish the same thing but encrypt traffic so that it is not easily readable.

Disable HTTP web access to web servers built into Ethernet attached products, and use HTTPS instead.

HTTP is the web access version of Telnet. HTTPS is the encrypted version. You should not use HTTP for the same reasons you should not use Telnet.

If you really want to ratchet down security, you can disable Telnet, SSH, HTTP, and HTTPS and rely instead on the serial console port for direct connection to a computer. But be sure to verify that all configurations required can be made through the console port before disabling the above services.

Disable ports that are not in use.

If you have ports that are not being used on an Ethernet switch, disable them so that Ethernet devices plugged into these open ports will not be able to communicate over the network.

Enable MAC address filtering on Ethernet ports.

You should enable MAC address filtering for each Ethernet switch attached to your network. This means that ONLY specific MAC addresses can connect to specific ports on the switch. Unauthorized devices plugged into the switch will not be allowed to communicate.

Enable 802.1x authentication.

Before being allowed to communicate past an Ethernet switch, a credential check must be done with an authentication server (RADIUS and TACACS are common types). If you are not on the list of authorized users, you will not be able to get on the network.

Change default passwords; use best security practice passwords.

Never leave the default user name and password on your device. Default user names and passwords are just an Internet search away from being discovered.

Summary

Industrial networks have a different set of requirements than a commercial/enterprise network. A successful network deployment involves taking the time to analyze the requirements of the system which includes: QoS, cost of ownership and downtime, size and scope, and security. When the known requirements are established, a network can be engineered with redundancy, network segmentation, traffic prioritization, management systems, and security. You will not find a one size fits all formula, as each network has its own set of requirements.

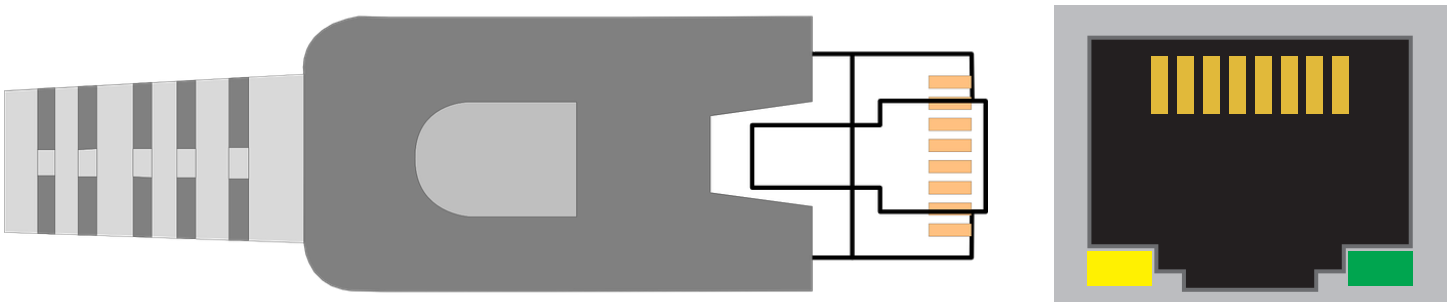

At Moxa, we understand that defining network requirements and design is not an easy task. To discuss your application with one of our experts, contact us at usa@moxa.com or +1-714-528-6777. For a Moxa office local to you, please visit www.moxa.com/about/Contact_Moxa.aspx.

About Moxa

With over 30 years of experience in communications technology for industrial automation, Moxa is one of the world's leading providers of solutions that enable connectivity for the Industrial Internet of Things.

Our edge connectivity, industrial computing, and network infrastructure products have helped connect more than 40 million devices worldwide in industries including: factory automation, smart rail, smart grid, intelligent transportation, oil and gas, and marine.

We pride ourselves in helping our customers harness the power of automation network convergence and making their operations smarter, safer, and more efficient.

As an active member of multiple industrial associations including the Industrial Internet Consortium, SunSpec Alliance, EtherCAT Technology Group, ODVA, Modbus IDA, and PROFINET International, we strive to promote open standards and interoperability.

Visit **moxa.com** to learn more about us and the different communications solutions that are tailored to the needs of industrial users, system integrators, and OEMs.

Follow us on social media!