

Securing EtherNet/IP™ Networks



Table of Contents

1	Introduction	3
2	Basic Overview of Risk and Security Approach.....	4
2.1	Risk and Risk Analysis	4
2.2	Reducing Risk.....	5
2.3	Costs and Tradeoffs.....	5
2.4	Security – Differences between IT and Industrial	6
2.5	Security – Working with IT	7
3	Best Practices for Different Types of Industrial Networks	7
3.1	Isolated Control Network with a Single Controller	7
3.1.1	General considerations	7
3.1.2	Managed Switches	8
3.1.3	Device Maintenance	9
3.1.4	End-Device Security	9
3.1.5	Network and Application Management.....	10
3.2	Isolated Network with Multiple Controllers	10
3.2.1	General Considerations	10
3.2.2	VLANs	11
3.2.3	QoS	12
3.2.4	Network Resiliency.....	12
3.2.5	IGMP.....	13
3.3	Enterprise-Connected and Integrated Control Systems	13
3.3.1	Firewall and DMZ	15
3.3.2	Intrusion detection.....	16
3.3.3	Remote Access	16
3.4	Wireless	17
4	Emerging Industrial Security Technologies.....	19
4.1	Security-Enhanced Operating Systems.....	19
4.2	IEEE 802.1X.....	19
4.3	Encryption and Virtual Private Networks (VPNs).....	19
4.4	Biometrics	20
4.5	Update throttling	20
4.6	NAC/NAP	20
5	Terms and Definitions.....	21
6	Additional Information	22

1 Introduction

Security is an essential element of network design and management in today's industrial enterprise. This guideline introduces the concept of "cyber-security" for EtherNet/IP™ networks and provides direction regarding important considerations for cyber security in industrial automation applications. The audience for this guideline is end users of EtherNet/IP who need to acquire a basic understanding of cyber security for industrial automation applications and to formulate an approach to achieve effective security practices. The document should be read in its entirety in order to gain a basic understanding of how to secure industrial networks that use EtherNet/IP technology.

In the past, dedicated automation control systems were unconnected beyond their specific application. Security was sustained simply by controlling physical access to the automation components. Today, however, connectivity to all the processes of the enterprise has increased productivity while reducing the actual time to market for new offers. This advance in connectivity has created a new path for both desirable and undesirable connections. This guide aims to outline issues to consider when deciding whether to bridge the automation network to the enterprise network and, eventually, the Internet. Security is discussed for the simple, stand-alone machine or process which is unconnected to the automation or company network, all the way through to the fully integrated network which links all the processes within an enterprise. Each "more-connected" example builds on the previous to provide a continuous path of value for the stakeholders.

Industrial networks have historically used fieldbus technologies that were not connected to other networks. However, users are moving to networks based on standard Ethernet and the IP protocol suite, most frequently to EtherNet/IP™, to connect these networks to the enterprise network and the Internet. This change provides a number of benefits, including increased visibility of plant floor activities, integration with back-office applications, and lower total cost of ownership. However, users need to be aware of how this connection impacts the security and availability of their industrial network and the automation and control systems they interconnect.

Security should be applied to any action that could disrupt normal business activity. This includes the loss of assets (including product, plant, production, or intellectual property), injury, and/or damage to personnel, products, tools, machines, the environment or company reputation. These disruptions can occur as a result of internal or external actions and could be intentional or accidental.

There is no "one-size-fits-all" solution to improving security. It is a complicated, multifaceted challenge that cannot be solved by simply purchasing the latest technology. Instead, managing the security of your industrial network requires changing processes and managing risk.

The practices described here are categorized by the different industrial network installations as described in ODVA's publication *Network Infrastructure for EtherNet/IP: Introduction and Considerations* (See Section 6 for information on where to obtain this publication.) Readers are encouraged to review this publication in its entirety, but should pay particular attention to the section that applies to their particular application. This document on cyber security assumes that the best practices recommended in ODVA's publication *Network Infrastructure for EtherNet/IP* will be applied to each different network installation. For example, there are security benefits of enabling IGMP Snooping, but it should be enabled as described in the EtherNet/IP infrastructure publication to optimize system performance.

Section 3 of this document provides information on the best practices for different types of network installations. The reader should read and understand the risks, tradeoffs, and best practices for their network application. If you plan on implementing an enterprise-connected industrial network installation, you should first read the previous section on isolated industrial networks with single and multiple controllers. For further description of the network types, reference *Network Infrastructure for EtherNet/IP: Introduction and Considerations*, in particular Ch. 6 - Infrastructure Application Scenarios.

Section 6 of this guideline includes a list of recommended reference documents on various topics. This document is not meant as a detailed explanation or checklist on how to secure an EtherNet/IP network but rather as a starting point in the user's education and investigation. Topics of precautions and mitigation steps related to industrial automation, along with lists of potential security risks, are beyond the scope of this guideline but reference documents are included in Section 6.

2 Basic Overview of Risk and Security Approach

The first step in determining a security strategy for your EtherNet/IP network entails identifying your potential risks in order to minimize them.

This section highlights some key steps required to develop a security approach for automation and control systems, in particular identifying and analyzing risk and taking steps to reduce that risk. These concepts are expounded upon in ISA 99's "**Security for Industrial Automation and Control Systems: Establishing an Industrial Automation and Control Systems Security Program**" documents (ANSI/ISA-99.02.01-2009). An earlier document from ISA99 is also useful in this context: "**Security for Industrial Automation and Control Systems Part 1: Terminology, Concepts, and Models**" (ANSI/ISA-99.00.01-2007). Best practices described later in this document may not necessarily be highlighted in the ISA99 documentation, but it is not the scope of the ISA99 document to highlight the best practices of securing automation and control systems.

2.1 Risk and Risk Analysis

Before discussing the specific cyber security risks in industrial networks, it is important to understand the basic concepts of risk and risk analysis. The risk for any particular device or system is the expectation of loss expressed as the probability that a particular threat will exploit a particular vulnerability with a particular consequence¹. While security incidents in the information technology (IT) environment usually relate to the loss or corruption of information, in the industrial environment, the consequences of cyber security incidents can have physical effects on production or the health, safety, or environment of the organization and the surrounding community.

When industrial networks were isolated from other networks, the risk was managed by ensuring that the plant was physically secure. The common saying that a user can do more damage with a hammer was true: if the user wasn't physically on the premises, they could not inflict damage.

However, once users started adding dial-up modems for remote troubleshooting, the plant was open to attack. As users connect or integrate their industrial network with other

¹ ISO/IEC 15408-1: Information technology – Security techniques – Evaluation criteria for IT security: Security Architecture – Part 1: Introduction and General Model, Clause 4 (General Model)

networks, including the enterprise network, the risk of a security incident on the industrial network increases.

For each risk such as, Denial of Service or Illegal access, the following questions need to be asked:

- What are the consequences?
- What is the likelihood of the risk occurring?
- What is the cost to prevent vs. the cost of the impact?

When considering the likelihood, the user must consider both intentional and unintentional forms of security issues. For example, the likelihood of a virus-based attack is much higher than a hacking attempt directed at an average plant.

2.2 Reducing Risk

There are a number of general ways that users can reduce risk. The first way to reduce risk is to use a “Defense-in-Depth” approach towards security. There is no single device or method that will secure a network, so it is necessary to build a system that works together with multiple layers of protection. Defense-in-Depth applies to both the physical and electronic security of the network. To physically secure the network, access to the devices on the network should be controlled. In the same way that there are many people who should not be permitted on the factory floor, there are many people that should not have access to the industrial network. In addition, there are very few users on the factory floor that need access to all industrial applications – access to these tools should be limited as much as possible, without preventing users from safely getting their jobs done.

To electronically secure the network, multiple barriers or virtual walls should be installed around and within the network. This makes an attack more difficult in the first place and limits its spread if one does occur. This is explained in more detail in the description of each network type below, but an example of this in the real world is having a fence around a facility, as well as locked doors. If someone gets past the fence, he still needs a key.

The next way to reduce risk is to have a process in place to ensure that all devices have the most recent security patches and anti-virus updates. A number of security incidents occur when someone is able to take advantage of a known security flaw or virus (reference section 3.31 for additional information on security appliances).

Another way to reduce risk is to minimize time to recovery. Regardless how many different ways users prevent their system from being attacked, users should be prepared to handle such an incident. Users should have copies of system configurations, plant diagrams, etc., stored in a secure location for disaster recovery purposes.

2.3 Costs and Tradeoffs

There are significant benefits to connecting automation and control networks with enterprise networks versus leaving the automation and control isolated. But there are tradeoffs that must be made between risks and costs. Security is about minimizing the risks and threats while taking maximum advantage of the benefits. For example, in connecting the plant to the enterprise, one may allow certain types of traffic flows and applications to communicate, but restrict others that have more risk involved.

EtherNet/IP, like most industrial protocols, uses unencrypted messaging. The encrypting/decrypting of messages would significantly increase both the cost and processing delays in the end devices. Additionally, most automation and control networks are protected in a way thru network isolation (air-gaps) or thru numerous security techniques (Defense-in-Depth). For example, one method of increasing availability is to limit the EtherNet/IP traffic flow to the automation and control network, as mentioned later in this document. Limiting the flow of EtherNet/IP to trusted devices on the plant floor significantly reduces risk. This guide reviews a number of mechanisms to improve confidentiality, authenticity, and integrity in order to make the EtherNet/IP protocol and the industrial automation and control systems more secure.

Confidentiality, authentication, and integrity are normal parts of any secure communications over the Internet, such as when purchasing a book from an online retailer. Here the browser will automatically apply SSL (Secure Socket Layer), a secure set of protocols, to the normally insecure HTTP protocol as indicated on your browser by the **https://** in the URL and by the lock or key symbol in status bar at the bottom of the window. However, when you submit a form to a search engine your messages are sent unencrypted. There are fairly large and expensive communication devices at work in this process. The online retailer's server and your computer are both involved in the encryption process. Requiring EtherNet/IP devices to perform these encryption activities would either drastically slow down their communication rates, slow down their ability to perform their control functions, or require very expensive CPUs to be installed in these devices. Delays due to encryption, decryption and increased CPU processing overhead simply cannot be tolerated in most automation and control systems.

Security techniques that can be applied today are discussed later in this guide. Because EtherNet/IP adapts commercial-of-the shelf technology used in standard unmodified Ethernet, ODVA will continue to improve the security features available in EtherNet/IP as the costs of deploying these features decreases and/or their value increases.

2.4 Security – Differences between IT and Industrial

The IT and industrial departments have different methods to achieve their goals due to their different requirements. IT networks, outside of data centers and servers, have relatively low requirements for determinism and availability. A user can wait multiple seconds for a web page to load or can wait multiple hours for a problem to be fixed. Industrial networks, on the other hand, have relatively strict requirements for both determinism and availability. Many industrial processes require message timings on the order of tens of milliseconds and 99.999% availability. In the future, the determinism and availability requirements for both groups may become more stringent as the IT department adds voice and video traffic on their network and as multi-axis motion control and safety is added to industrial networks.

The IT department achieves its security goals by providing multiple layers of security. For example, one layer of protection is provided by firewalls that separate the entire enterprise network from the Internet and other enterprise networks. These firewalls inspect all incoming and outgoing packets, and drop any packets that might be harmful. Within the enterprise network, another layer of security is provided by placing limitations on who can access a set of data. (For example, Human Resources might be the only group that can make changes to the Human Resources database that exists on a server.) Another layer of security is provided by requiring all servers and PCs on the network to have the latest anti-virus and OS patches.

The IT department applies these policies to maintain uptime for the maximum number of users and systems. As a result, it is permissible for a single system to be down to receive an anti-virus or OS patch because it prevents downtime for a large number of users. However, for servers that provide service to a large number of users, anti-virus and OS upgrades are typically applied only during scheduled downtime.

Industrial applications have different requirements compared to IT requirements, and therefore require different methods to secure the network. While enterprise networks will continue operating if a single device goes offline, in most cases, the production process will stop if a single device goes offline or is processing non-production information. For example, running an anti-virus scan might prevent a device on the industrial network from responding to commands.

2.5 Security – Working with IT

In many cases, the IT department might try to apply the same processes and tools to automation systems that are placed on desktop machines – for example, an OS patch is applied the moment it becomes available, and the machine is forced to reload. This is clearly an unworkable solution for automation systems.

Instead, the automation department should work with the IT department to explain the automation requirements – in many cases, these will match the same requirements for data center systems. In a data center, patches are applied during scheduled downtime to prevent any impact to the production network – the same thinking/process can be applied to automation systems.

It is also important for the automation department to explain the traffic patterns of their equipment to the IT department. In most cases, devices that communicate via EtherNet/IP do not access or send information with devices outside of the company. As a result, the risk of a virus infiltrating into the enterprise network from an EtherNet/IP device is very low. Also, based on traffic patterns, network filters and firewalls can be configured to prevent security problems on the enterprise network from affecting devices on the industrial network and vice versa.

3 Best Practices for Different Types of Industrial Network Installations for EtherNet/IP

This section will break down the best practices by the types of industrial network installations as identified in *Network Infrastructure for EtherNet/IP: Introduction and Considerations*. These network installation types represent the level of interconnectivity between the industrial and enterprise networks. The intent is that the security approach should align with the size and connectivity of the network. It is also understood that implementations of EtherNet/IP may develop and migrate with time and that the security considerations for the industrial network would parallel these developments. The best practices outlined with each type of network are therefore considered additive in nature. For example, the security best practices for an isolated control network with a single controller would also apply to an isolated control network with multiple controllers.

3.1 Isolated Control Network with a Single Controller

3.1.1 General considerations

Of all the categories of industrial networks, this category is the smallest and least complex as shown in Figure 1. These may be in small, single-operator shops or there could be a

large number of isolated work cells within an organization. While they only have a single controller, they may have a large number of adapters and I/O points that require multiple switches.




Figure 1 - Example Isolated Industrial network with Single Controller

Because these systems are considered small and isolated from the enterprise network, the risks associated with this type of network are limited. An attacker (whether unintentional or intentional) would have to be in direct contact with the network to affect its operation. The main threat would be from infected computer resources such as laptops, USB memory sticks and other media attached to the network or a computer on the network. Users should scan all devices prior to connecting them to this system, or have a company-owned secure laptop available for users that need to connect for maintenance or debugging purposes. Even systems that do not have a virus can impact availability if, for example, they are configured to act as a DHCP server or have incorrectly configured network settings.

Another possible threat includes the destruction or manipulation of the controller code (either unintentional or intentional). Since these systems usually don't have multiple operators, there may not be any tracking of changes made to the controller or other devices on the network. The consequences associated with an incident on this type of network will usually result in the lack of availability of the controller or other resource. An incident may have health, safety, or environmental effects, but they will typically be limited to the location around the industrial network. Configurations should be backed up and stored in a secure location.

3.1.2 Managed Switches

While not required for performance reasons in an isolated control network, managed switches can improve the security of the network. Managed switches can be configured to limit the traffic rate on a per-port basis, using known traffic patterns, via port-based security (e.g., MAC or IP port security). The management features (e.g., QoS, IGMP) of these switches can also improve the security of the network. The impact of a network storm as a result of a virus or damaged equipment can be minimized by using this functionality. Users should be careful to configure traffic filters so that normal traffic isn't blocked. Switch ports that are not regularly used should also be disabled to prevent accidental use/connection to the industrial network.

3.1.3 Device Maintenance

It is fairly common for larger users to have maintenance contracts on some of their devices that require a technician to monitor and perform maintenance on a device regularly. This device maintenance can either be done locally at the site of the device or via secure remote access from the maintenance provider's site.

If the maintenance is being conducted locally, a strict policy must be enforced on access to the industrial and enterprise networks. Individuals should not be allowed to connect unknown devices (computers, network equipment, etc.) to the industrial network without first being checked for current anti-virus updates, software patches or compatibility with the network and applications.

If the maintenance needs to be conducted remotely by any means (dial-up phone line, cellular router, VPN, Internet, etc.), then the network is considered to be enterprise connected and integrated system (see sections 3.3 enterprise-connected and integrated control systems, below) and should be treated as such. A security policy and procedure must be enforced, dictating the authorized users and activity for this connection. The remote connection should use a network segmentation device and should be monitored for any activity outside the recognized security policy.

If the device has a web interface or SNMP, it is recommended to change the default password and avoid posting the password in a public or non-secure location. It is also a good practice to disable any un-necessary ports and services.

3.1.4 End-Device Security

Devices in the industrial network running a common operating system (OS) provide an opportunity to introduce malware such as a virus, worm, Trojan horse or other common end-device attack which usually target the common operating systems. Anti-virus software, which usually protects against a range of known malware attacks, and regular patching are all common mechanisms to reduce the potential of an attack or downtime due to these types of influence. Naturally, an end-device in an industrial network may not be able to be patched as easily or regularly as, for example, an enterprise computer, but a regular maintenance schedule should be developed and kept. Many embedded systems such as the PLC/PAC or EtherNet/IP remote I/O do not use these operating systems. These embedded systems are less complex and do not support as many networking features as an office PC, and therefore require fewer security updates.

Since many EtherNet/IP devices use non-IT hardware and operating systems, the number of viruses, worms, Trojan horses, etc. has remained low or non-existent. However, industrial automation and control systems may be impacted by standard denial of service attacks on the network or EtherNet/IP PC-based devices could be affected through standard email, webpage, and file exchange attack methods.

End devices with common operating systems such as a Windows-based machine should have security applied for protection such as virus software and should be upgraded and maintained on a regularly scheduled basis. Additionally, the use of browser and other internet applications has been a significant source of security breaches and attacks. Consideration to limit the internet access or network accessibility of end-devices in the production environment is another important security consideration.

3.1.5 Network and Application Management

Network and application management play a key role in any automation and control security approach. Although they may not stop an attack before it happens, monitoring network and application services is key to recognizing and reacting to attacks or breaches.

For example, attacks based on sending malformed packets have been shown to allow an attacker to either disrupt or take over commercial or industrial devices. Malformed packet attacks are possible due to incomplete or non-robust implementations of the existing TCP/IP suite and industrial protocols. Either way, intentional or unintentional, malformed packets and other improperly constructed communication can negatively impact performance or may be used to breach a device.

Managing and monitoring of the network and automation and control devices for such things as CIP errors will help identify and stop these threats or at least identify possible security breaches. Best practices include setting thresholds in the end devices and controllers to warn operations personnel that abnormally high packet failures or other unexpected conditions have occurred. Similarly, monitoring and management of key network statistics and errors can help prevent attacks targeted at both end devices and the network infrastructure itself.

As well, use of encryption for access to the network infrastructure is an IT best practice suited for plant networks. Use of SNMP v3, SSH and HTTPS for accessing and managing infrastructure devices is included among these. Encryption is also accompanied by use of authentication and authorization for access to network infrastructure (logins, passwords and access to individual parameters). As well, simple actions like posting banners on login pages to indicate the type of switch being accessed can help limit errors or unintentional mistakes.

3.2 Isolated Network with Multiple Controllers

3.2.1 General Considerations

Larger installations may need more than one controller on the industrial network, but corporate policy may require that the industrial network be isolated from the enterprise network as shown in Figure 2. Networks of this type can have a multi-layered architecture using managed switches and Virtual Local Area Networks (VLANs) to segment the larger number of devices including local servers. Controllers can be put into different VLANs to improve overall system performance and availability by separating traffic between devices.

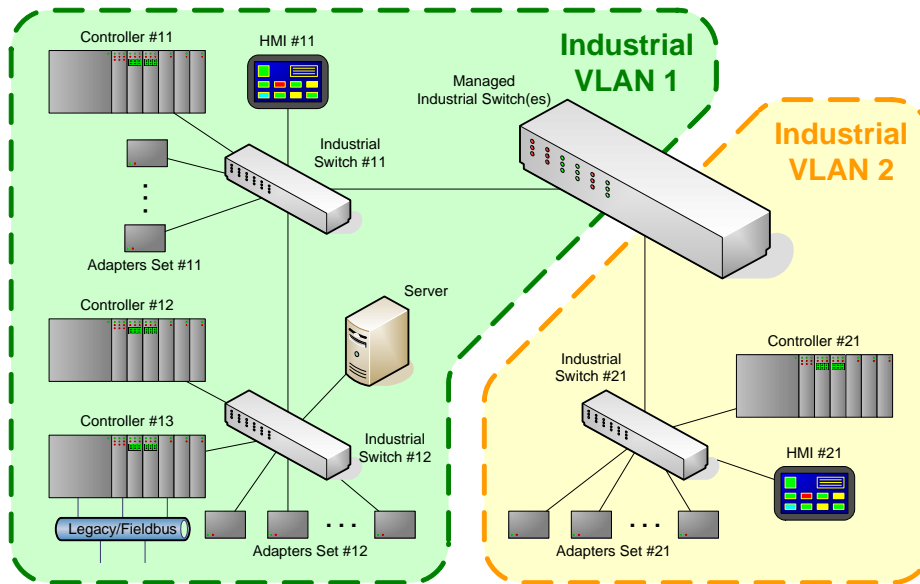


Figure 2 - Example Isolated Industrial Network with Multiple Controllers

Due to the additional complexity of these types of networks, it is possible for an incident at one end of a facility to affect a device on the other end of the facility, although VLANs help limit that impact. An attacker (whether unintentional or intentional) would still need to be in direct contact with the network to effect plant operations; however, this may occur more often than in the single controller network described above. A larger organization may have contractors that work alongside employees maintaining or operating equipment. The industrial network may exist throughout an entire facility and may have many open and susceptible network ports.

As with the single controller network, one of the main threats to the network would come in the form of infected computer resources. Another threat that is shared between single and multiple controller systems would be the unintentional or intentional destruction or manipulation of the controller code. In the multiple controller case, it may be due to an attacker maliciously attempting to affect the process or it could come from the plant engineer uploading a program to the wrong controller. In general, these incidents would be very similar to the single controller case, but the consequences would not be limited to affecting only the local area. An incident at one location may affect the area surrounding the controller, another area at the facility, or the entire facility. Incidents on this type of network would not typically affect multiple facilities unless a dedicated industrial network had been configured between those facilities.

Networks with multiple controllers can easily be larger and more complex than any one person can manage. Due to this fact, it is a good practice to develop more detailed policies and procedures to ensure security practices are being followed by personnel. The industrial network should be designed to protect the devices and controllers from inadvertent events that may disrupt normal operations.

3.2.2 VLANs

VLANs are a method of creating logically-independent networks within a huge network interconnected through switches. A VLAN consists of a network of computers that behave as if connected to the same wire - even though they may actually be physically connected to different segments of a LAN. Network administrators configure VLANs through software

rather than hardware, which makes them extremely flexible. Several VLANs can co-exist within such a network. This helps in reducing the broadcast domain and therefore the effects of broadcast storms, worms and viruses, and Denial of Service (DoS) attacks. VLANs are therefore often deployed to improve the overall system performance and availability, which are important security considerations. But, VLANs also bring significant security benefits in addition to performance and availability. They form the basis for a security policy by segmenting the devices. For example, security policy and implementations may be based on VLANs rather than individual devices, thereby simplifying the effort to implement and maintain a security approach.

3.2.3 QoS

Quality of Service, or QoS, features can be used to ensure that high-priority traffic isn't affected when a network storm occurs. Without QoS, all packets have the same priority as they are sent across the network. When there is a network storm for any reason, EtherNet/IP traffic must compete with the network storm. When this happens, the EtherNet/IP traffic might have to wait behind other packets, or might be dropped. When QoS is enabled, high-priority packets are sent before low-priority packets, and low priority packets are dropped before high-priority packets. In many cases, enabling QoS will prevent a network storm from affecting the industrial network. Therefore, QoS is often considered a best practice for a security policy for industrial Ethernet networks. QoS is also a best practice for highly available or congested networks.

ODVA has now included QoS guidelines in *The EtherNet/IP Specification*, which also requires that the low-latency CIP Motion and CIP Sync traffic producing end-devices mark its traffic so it receives higher priority as the packets traverse the network infrastructure.

3.2.4 Network Resiliency

EtherNet/IP networks must be highly reliable and continue to operate during harsh environmental conditions, accidental network disruptions, and equipment failures. Network downtime can be dangerous and expensive. A common aspect of network resiliency can be provided through redundancy, coupled with firmware in the device that instructs the network to switch to alternate paths upon specific failures. In other words, resiliency can be achieved by forming a backup path when part of the network becomes unavailable. One of the first technologies developed for this application was the open standard IEEE 802.1D Spanning Tree Protocol (STP). Although IEEE 802.1D STP has solved some resiliency requirements, it also has limitations including slower convergence speed, constraints of bridge diameter, VLAN insensitivity, and link blockage (when the bandwidth is not enough for all traffic). For this reason, IEEE 802.1W Rapid Spanning Tree Protocol (RSTP) was developed. This newer protocol has all the advantages of IEEE 802.1D, but in addition provides higher performance. For industrial applications requiring the fastest communication recovery times, proprietary ring protocols are commonly used. Consequently, many industrial Ethernet switch manufacturers have developed proprietary protocols, some based on 802.1W enhancements to meet the most stringent recovery times required in industrial automation. Proprietary ring protocols can often provide recovery times in the millisecond range for networks with over 100 devices. However, the proprietary nature of the protocols often requires that all switches are manufactured by the same vendor in order to maintain those aggressive recovery times. If the network topology requires a mixing/matching of switch vendors then 802.1W (RSTP) becomes the most suitable default.

Device Level Ring (known as DLR) (IEC 61158-4-2: 2010(E), clause 10) is a ring topology included in *The EtherNet/IP Specification*. DLR can be used as a means to quickly converge a ring of automation and control devices. DLR is a protocol intended for end-devices

containing an embedded switch with two external ports. The quick convergence time, under 5ms for rings of 50 devices (nodes), means that the ring nodes can detect a break in the ring and reconfigure the network fast enough to avoid the loss of I/O connections. This allows the automation and control application to continue to operate normally, without interruption.

3.2.5 IGMP

The Internet Group Management protocol (IGMP) provides a mechanism for the networking infrastructure to pass multicast traffic to only the end-devices that request the information. IGMP is considered a best practice in many EtherNet/IP systems to provide multicast management support as much of the critical data transfer between controllers and I/O is done via multicast messaging. The multicast management ensures that the end device receives only the messages they need and that the network only has to process wanted multicast traffic, thus improving the performance of both end devices and the network infrastructure.

Other security benefits of IGMP are that devices only receive the multicast packets that the network infrastructure understands the end device wants.

3.3 Enterprise-Connected and Integrated Control Systems

Some organizations may choose to connect their industrial network and enterprise network, as shown in Figure 3, for added flexibility and functionality. If the industrial network could not operate properly without the enterprise network or requires real-time traffic to be passed between industrial networks, then it is termed as an enterprise-integrated industrial network. The amount of data passing over the enterprise infrastructure is not important in order to consider a network enterprise integrated. The industrial network(s) may pass low-frequency data like inventory statistics to an ordering system or high-frequency data like real-time I/O data between work cells on a plant floor, both are considered enterprise connected.




Figure 3 - Example Enterprise Integrated Industrial network

The risk of attack to the industrial network increases greatly when it is attached to the corporate infrastructure in any way. This risk can be reduced if both networks are designed with security in mind and appropriate segmentation is applied. Enterprise-connected industrial networks have all the same threats and vulnerabilities that isolated industrial networks have in addition to many other associated with the enterprise connection.

Many of the recent computer incidents that organizations have faced have come in the form of email- or webpage-related viruses, worms, or Trojan horses. Once the industrial network is attached to the corporate infrastructure, there is also most likely access to the Internet. Any employee or contractor in the organization, regardless of their location, could affect the operation of the industrial network. An incident on the industrial network could affect the enterprise network, or vice-versa.

Since the industrial networks may not operate properly without being interconnected through the enterprise network, a security incident in any part of an organization (including other facilities) could have an effect on the industrial network. The attack would not have to be directed at the industrial network, it might only be directed at the enterprise network, but that would impact the operation of the connected industrial network(s) as well.

In addition to the internet and data concerns, connecting the enterprise to the industrial network also opens up the possibility that other types of network traffic may maliciously or unintentionally traverse that connection and affect production and control.

This type of connectivity is not new and is used successfully in a many networks, and, if achieved in a secure manner, can offer compelling benefits that:

- Include tighter and more real-time access between enterprise and industrial services and data.
- Include secure Remote Access to allow experts and plant personnel access to plant networks and devices from remote and potentially external locations, greatly reducing the need to travel or have additional experts on-hand to quickly resolve problems and issues that arise.
- Enable plant personnel or guests access to enterprise or external services and information to make them more productive while connected to the industrial networks or located on-site.

Some of the best practice precautions include:

- Using network infrastructure to highly segment the enterprise and industrial networks, for example with firewalls.
- Develop a "De-Militarized Zone" (DMZ) where data and services can be shared, but is separate from the industrial or enterprise zone.
- Develop techniques and services to provide Remote Access.
- Replicate or place critical network-based services into the industrial networks (e.g. active directory servers, DHCP servers, print servers, etc.).

Of particular interest in securing an enterprise-integrated industrial network is the configuration and use of network switches. Managed network switches have numerous

options that can be used to address security and performance. Some of the options have default settings and others need to be activated and configured or tuned to improve security. The complexity and effort involved in utilizing many of these options must be traded off with the perceived risks and costs. The options most likely to affect control networks are considered here.

Layer 3 switches, network routers and firewalls have the ability to control which devices can communicate through specific port combinations of the switch. This can be used to restrict the sources and destinations of traffic between the enterprise and the industrial networks. A common use of this capability would be to only allow plant, engineering, or maintenance managers access to the industrial network from their enterprise connected office computers.

An enterprise-integrated industrial network should be expected to be under continual change and upgrades, especially on the enterprise side. Because of this dynamically changing environment, it's good policy to periodically perform security checks or audits to reduce the risk of changes causing a reduction in security or an introduction of additional threats.

If these networks are connected through a CIP gateway, regardless of network size and complexity, it is called an enterprise-connected industrial network. All traffic between the industrial and enterprise networks would be filtered through this CIP gateway, limiting the traffic to EtherNet/IP-only messages. Since the CIP gateway would be required to filter each incoming/outgoing message, this type of network is not recommended for situations that require high-speed, real-time interaction between the industrial and enterprise networks. The CIP gateway should not be required for production, since a traffic storm could affect its normal operations.

Connecting the enterprise network to the industrial network introduces the set of risks associated with transporting and sharing data that is not critical to the operation of the industrial network. Of concern might be the confidentiality of the shared data as it traverses the enterprise and who has been authenticated and authorized to access such data.

Many managed switches also have the ability to examine the traffic they route at the bit, byte, and/or packet header level. This gives the switch the ability to make a judgment and decision on how to handle current and/or future traffic. For example, a switch could be configured to permit all EtherNet/IP traffic but block e-mail traffic. Configuring a switch to strictly limit the types of traffic permitted has the benefits of additional security, but can be difficult to manage as new applications are added.

3.3.1 Firewall and DMZ

In building construction, a firewall is a wall designed to stop the spread of fires and other threats within the building. Buildings are often segmented into zones with firewalls to contain threats such as fires to a zone and preserve the other zones, at least for a period of time. Much the same applies for network firewalls. They are designed to provide strict segmentation of network traffic passing between two or more zones and restrict or stop the spread of attacks.

A common and efficient means to restrict and control the flow of information between the enterprise and industrial networks is through the use of a firewall. Network firewalls are already commonplace between enterprise networks and the Internet, and even between home networks and the Internet. A firewall between the enterprise and industrial networks

should be configured to only allow predetermined traffic to pass through. This might include allowing certain office personnel access to web pages in the control network that provide production, supply, and maintenance data. Another possibility would be to allow devices on the control network to push (e.g., via email or FTP) their appropriate data out to predefined enterprise devices.

Firewalls are limited though, and many lack the ability to perform packet inspection in applications dominated by feature and function-rich protocols such as the Common Industrial Protocol (CIP™), which is used by EtherNet/IP. Further, CIP has been optimized for functions needed in dedicated industrial automation applications and currently does support the authentication or encryption concepts common in internet computing. Therefore, many IT departments limit or block the movement of messages using industrial automation protocols such as CIP, through the firewall.

An additional consideration with a firewall is the creation of a network zone, the “De-Militarized Zone” (DMZ), used to share data and services between the enterprise and industrial networks. Traffic is allowed (although maybe limited) between the enterprise and the DMZ as well as between the Internet and DMZ, but not directly between the enterprise and the Internet. This is a common IT practice to share data and services from an enterprise with entities or users from the Internet and can be implemented between the enterprise and industrial networks. Of course, this entails the use of firewalls to manage the traffic flows and additional servers and network infrastructure to handle the data and services maintained in the DMZ.

Both firewalls and DMZ concepts are referred to in a growing number of industrial standards and guidelines, including NIST and ISA 99.

3.3.2 Intrusion detection

This is the effort involved in trying to detect actions that attempt to compromise the confidentiality, integrity or availability of the network or a resource.

Intrusion detection can be performed manually or automatically. Manual intrusion detection might take place by examining log files or other evidence for signs of intrusions. A system that performs automated intrusion detection is called an Intrusion Detection System (IDS). An IDS can be either host-based, if it monitors system calls or logs, or network-based if it monitors the flow of network packets. Modern intrusion detection systems are usually a combination of these approaches. Another important distinction is between systems that identify patterns of traffic or application data presumed to be malicious (misuse detection systems), and systems that compare activities against a ‘normal’ baseline (anomaly detection systems).

IDSs do not in general include prevention of intrusions, but can sometimes be coupled with subsystems that are designed to locate the source of the intrusion and temporarily or permanently block that source from attempting further network intrusions. IDSs can be complex and are best utilized by IT organizations at this time. At this point in time there are no IDSs that have rules designed to specifically detect EtherNet/IP intrusions. If firewalls are in place to block CIP traffic from passing between the enterprise and industrial zones, it is more difficult to intrude on EtherNet/IP communication.

3.3.3 Remote Access

Remote access is the idea that remote plant personnel or external experts can access a portion of the industrial automation and control systems in a plant. Remote access can

suggest access from either within the enterprise or, external from the internet. Secure remote Access of industrial networks is a key benefit of a strong security design. For a variety of reasons, manufacturers are always struggling to find experts to help deploy and maintain the industrial automation and control systems and the industrial networks. Remote Access offers a means to more flexibly and yet securely bring that expertise into the plant.

To offer secure remote access, many of the previously mentioned security best practices are recommended, including:

- Firewalls/DMZ provide a key choke point for network traffic and allow tight control of the access into the industrial network, including authentication enforcement, limiting the type of applications that can be used to access the plant applications, proxy to hide details about the network and devices, and if, enabled, IPS to monitor the traffic for known security attacks.
- Terminal services come in a variety of forms now, either direct, such as Citrix, Remote Desktop Services (MS based) or Virtual Network Computing (VNC), or via web conferencing tools. The DMZ enables the service to be hosted in a network zone where both external and internal personnel can access, but direct communication may be limited.
- VLANs for segmentation to limit the remote personnel's ability to access networked devices.

In addition, key considerations for remote access include:

- Control the software used. Plant personnel should have a defined set of approved software tools.
- Control or limit the capabilities of remote personnel. Often, line-of-site or other location restrictions are required and need to be maintained while in the remote access mode.
- Monitor and audit the activities of the remote personnel, so if mis-use occurs, it can be tracked and identified.

3.4 Wireless

Wireless transport of data is becoming much more common in industrial automation applications. The use of Wireless LAN based (i.e., IEEE 802.11) technologies is especially straightforward between and within installations using EtherNet/IP due to the use of TCP/IP protocol by EtherNet/IP. Other wireless technologies may also be deployed, but are not considered in this section.

Users with wireless access points should read the section on Enterprise-Connected Industrial Networks as a wireless network requires similar security considerations. But wireless technologies have additional and separate security considerations due to the use of a common medium: airspace and radio frequencies.

An emerging threat to control networks is the ubiquitous laptop that has built-in wired and wireless network connection capabilities. Such laptops have the ability to connect simultaneously to both a wired and wireless network. While these laptops do not by default

act as a medium between two such networks, that functionality is extremely easy to configure. This is another good reason to only connect company-owned, secure laptops to the industrial network.

Industrial wireless communications are becoming increasingly popular. The convenience of sending Ethernet packets over the air instead of through a wire can lead to many benefits. The IEEE 802.11 standard (commonly referred to as Wi-Fi) established a way to use radio frequency technology for Ethernet communications. Consequently these wireless local area networks (WLANs) which are prevalent in numerous commercial applications are being implemented in plant floor environments as well. However, some inherent issues which present a challenge include security, reliability, throughput and latency. Initial security technologies included WEP, WPA and now WPA2 (802.11i). WEP and WPA have known security limitations and should be avoided whenever possible. The use of WPA2 is considered a best practice. 802.1X in conjunction with wireless encryption can also improve the security of the wireless communication.

IEEE 802.1X is an IEEE standard for port-based network access control ("port" meaning a single point of attachment to the LAN infrastructure). It provides an authentication mechanism to devices wishing to connect to a LAN, either establishing a point-to-point connection or preventing it if authentication fails. It is used for most wireless access points. Access point vendors now use 802.1X to address the security vulnerabilities previously found in WEP. The authenticator role is either performed by the access point itself via a **pre-shared key** (referred to as [WPA2-PSK](#)) or for larger enterprises, by a network service, such as a RADIUS (Remote Authentication Dial In User Service) server. This provides for strong mutual authentication.

In general, effective authentication and encryption are the keys to a secure wireless network. Authentication is the positive identification of a network entity, such as a client or a server. Site authentication has been standard on secure servers for some time, because users require assurance that the data they receive from a site is actually being transmitted by that site, rather than by an imposter or eavesdropper. Encryption is the encoding of data in order to hide its content from everyone except its intended recipient.

Additional wireless security concepts include:

- **Rogue access point & device detection:** Rogue access point and device detection is a two step process starting with discovering the presence of an Access Point/wireless device in the network and then proceeding to identify whether it is a rogue or not. Rogue devices can potentially disrupt wireless networks and can sometimes cause irrevocable damage. Companies could unknowingly open up their Intellectual Property to outsiders and competitors through a poorly configured or unauthorized wireless device. Users deploying Wireless LANs should effectively detect and block wireless access points and client stations automatically and in real-time.
- **Disabling SSID broadcasting:** Most wireless access points automatically transmit their network name ([SSID](#)) into open air at regular intervals (every few seconds). This feature of [Wi-Fi](#) network protocols is intended to allow clients to dynamically discover and roam between [WLANs](#). However, this feature also makes it easier for hackers to break into a network. Since SSIDs are not encrypted, it becomes easy to grab one by snooping the WLAN looking for SSID broadcast messages coming from the access point. Knowing the SSID brings hackers one step closer to a successful intrusion.

- Key maintenance (WPA)
- Segmentation for enterprise or guest access

4 Emerging Industrial Security Technologies

As security threats expand and change, and computer and networking technology improves, security technologies must evolve to keep up. Here are some of the emerging technologies that may affect security for future control networks.

4.1 Security-Enhanced Operating Systems

While only a very small percentage of industrial networks currently contain devices using security-enhanced operating systems, there is a cost and security/reliability trend to move towards incorporating these operating systems into control networks. One example of a security-enhanced operating system is SELinux, which is an effort sponsored by the National Security Agency (NSA) to develop and deploy an even more secure version of Linux.

4.2 IEEE 802.1X

802.1X is an IEEE standard that allows for authentication of devices attached to a LAN port. Before a switch forwards packets through a port, the attached devices must be authenticated via an authentication server (commonly known as a RADIUS server). As a result, devices that are unknown to the network are not given access, or are put onto a VLAN that provides limited access. This is different from the username/password combination that is typed into most computers today. A benefit of 802.1X is the switches and/or access points themselves do not need to know how to authenticate the client. All they do is pass the authentication information between the client and the authentication server. If a device cannot complete the 802.1X process successfully, the device may be put on a remedial VLAN or even have its network port administratively turned off.

Currently, most industrial devices do not support 802.1X. 802.1X is a common technology in enterprise wireless networks and a growing feature in enterprise wired environments. It was highlighted earlier as a best practice for wireless industrial networks. For industrial applications, it may provide value for devices that can support an 802.1X client. 802.1X in industrial networks is not yet a common or best practice, but may become such in the future.

4.3 Encryption and Virtual Private Networks (VPNs)

Encryption provides a way to conceal information and make it unreadable to a third party. Encryption is commonly used when data is sent on an un-trusted network and/or if the data is highly confidential. There are a number of different methods to encrypt information, depending on the application.

Wireless networks should always be configured to use strong encryption. The type of encryption used will depend on the features that a device and/or wireless access point supports.

One type of encryption that is commonly used is a virtual private network, or VPN. A VPN uses cryptographic tunneling protocols to provide the intended confidentiality (blocking snooping and packet sniffing), sender authentication (blocking identity spoofing), and message integrity (blocking message alteration) to achieve privacy.

When using a wired network, packet encryption or a VPN is usually not required when traffic does not go beyond a single secure facility. Traffic sent between locations may need to be encrypted or sent over a VPN, depending on the sensitivity of traffic, knowledge of the networks the traffic will pass over, and latency requirements. EtherNet/IP traffic might not need to be secured if it is just being sent over the IT network. However, it might need to be secured if the traffic is being sent along an oil pipeline.

4.4 Biometrics

Biometrics provides the ability for a device to authenticate an individual based on a physical or behavioral attribute, such as a retinal scan, fingerprint, or voice recognition. Biometrics could be used to verify that a user is permitted to use a computer that configures devices on the industrial network. Utilizing biometrics for personal authentication is becoming convenient and considerably more accurate than prior methods (such as the utilization of passwords or Personal Identification Numbers). The main reasons are because biometrics links the event to a particular individual while a password or token may be used by someone other than the authorized user. Additionally, it is convenient (nothing to carry or remember), accurate (it provides for positive authentication), can provide an audit trail, and is becoming more cost effective.

4.5 Update throttling

Some operating systems (e.g., Linux) and applications (e.g., Peer-to-peer applications like BitTorrent) already have the ability to throttle network flow. A common complaint in industrial control networks is of the processing and networking overhead required by a computer to download and apply an operating system patch or anti-virus update. Disruptive loads of this type in the control system network can have disastrous effects. However, as more industrial nodes move to Ethernet and TCP/IP based networks, market pressures and advancing technology may bring about the capability to easily perform network, disk, and CPU rate limited patches and anti-virus and malware updates.

4.6 NAC/NAP

Network Access Control (NAC) and Network Access Protection are developing complementary technologies to validate the end-devices, their security stance and software levels, including for example the level of anti-virus updates and based on that, allow (or deny) access to the network. Network Admission Control (NAC) is a set of technologies that uses the network infrastructure to enforce security policy compliance based upon a number of characteristics of the end device. NAC is often considered as including user-base authentication mechanisms such as 802.1x. Network Access Protection is a Microsoft technology that implements NAC for Microsoft-based end-devices.

NAC and NAP offers a lot of potential for industrial networks as network access may be determined using other means than user-based authentication. Most industrial network devices (controllers, sensors, drives, etc.) do not have “users” on which to base network access decision. Other means, such as MAC address, IP Address, traffic patterns etc., may be used to validate network access. NAC and NAP offer the ability to validate every device on the network and monitor their behavior.

NAC and NAP are developing technologies. Due to the nature of allowing or denying access to the network, caution and sufficient testing should occur before implementing any of these technologies.

5 Terms and Definitions

- **Authentication:** A process to verify that the user or device is who they claim to be.
- **Authorization:** A process to verify that the user or device is allowed to perform an activity. Authorization can also be based on geography, i.e., User X in Location X can perform Activity X, but not when located in Location Y.
- **Availability:** The condition where devices are functioning properly without interference. Availability can be impacted by malfunctioning devices, security attacks, etc.
- **CIP Gateway:** A controller or device with two or more EtherNet/IP modules used to separate CIP traffic between different networks, such as the industrial network and the enterprise network. The gateway will block any non-CIP network traffic between the networks.
- **Confidentiality:** The ability to prevent unauthorized users or devices from reading data.
- **Integrity:** The ability to protect against unauthorized modification or destruction of hardware, software, or information in the system.
- **Malware:** Software used to attack a computer or device without the user's consent.
- **Quality of Service (QoS):** The process of sending and receiving packets with different levels of prioritization. When there is an overflow of traffic, low-priority packets are dropped before high-priority packets.
- **Virtual Local Area Network (VLAN):** A technology used to segment traffic. In order for traffic to pass between two different VLANs, the traffic must pass through a device like a router, firewall, or CIP Gateway.
- **Wireless Local Area Network (WLAN):** A local area network that uses radio frequency communications between devices instead of a physical cable media.

6 Additional Information

This document is meant to provide high-level information on securing industrial networks. The following links provide additional information on securing industrial networks, or were referenced in this guide.

- [1] **IAONA Security Handbook**
http://www.odva.org/Portals/0/Library/Publications_NotNumbered/IAONA_SysAsp_IPv4ACD_v1.0.pdf

- [2] **ISA-SP99, Manufacturing and Control Systems Security**
<http://www.isa.org/MSTemplate.cfm?MicrositeID=988&CommitteeID=6821>

- [3] **ODVA Network Infrastructure for EtherNet/IP: Introduction and Considerations**
http://www.odva.org/Portals/0/Library/Publications_Numbered/PUB00035R0_Infrastructure_Guide.pdf

- [4] **SANS Introduction to Learning about Network Security**
http://isc.sans.edu/presentations/first_things_first.html

- [5] **Using Host-based Anti-virus Software on Industrial Control Systems: Integration Guidance and a Test Methodology for Assessing Performance Impacts**
http://www.uscert.gov/control_systems/practices/pcsf/groups/d/1177076007-nist_sp1058.pdf

- [6] **Special Publication 800-82, Guide to Supervisory Control and Data Acquisition (SCADA) and Industrial Control Systems Security**
<http://csrc.nist.gov/publications/nistpubs/800-82/SP800-82-final.pdf>

- [7] **US-CERT Control Systems - Standards and References**
http://www.us-cert.gov/control_systems/csstandards.html

About ODVA

Founded in 1995, ODVA is a global trade association whose members are comprised of the world's leading automation companies that make and sell products compliant with ODVA technologies. ODVA's mission is to advance open, interoperable information and communication technologies in industrial automation and thus create value for our members, adopters, alliance partners and employees. ODVA's vision is to contribute to the sustainability and prosperity of our global community by transforming the model for information and communication technology in the industrial ecosystem. For future interoperability of production systems and the integration of the production systems with other systems, ODVA embraces the adoption of commercial-off-the-shelf (COTS) and standard, unmodified Internet and Ethernet technologies as a guiding principle wherever possible. This principal is exemplified by EtherNet/IP – the world's number one industrial Ethernet network. To learn more, visit www.odva.org.

CIP, CIP Sync, CIP Motion, and EtherNet/IP are trademarks of ODVA. Other trademarks are property of their respective owners.