2000 IEEE Industrial Applications Magazine – Best Paper of 2000

PROTECT THAT NETWORK: DESIGNING SECURE NETWORKS FOR INDUSTRIAL CONTROL

Eric J. Byres, P. Eng.
British Columbia Institute of Technology
3700 Willingdon Avenue, Burnaby, BC V5G 3H2
Email: ebyres@bcit.ca

Abstract— With the evolution of data communications in process control, network problems have taken on new importance to process engineers. The data connections from DCS and PLC systems to the plant network are vital to production, yet can be an invitation to problems. This paper looks at several real-life network disasters and discusses strategies for avoiding them. Solutions, including the use of packet filter firewalls and VLANs (Virtual Local Area Networks), are discussed. A case history from a pulp and paper mill illustrates how a firewall can be implemented to protect process systems from business users.

I. THE NEED FOR NETWORK PROTECTION

Over the past ten years the process control field has seen a significant increase in the use of computer networks to transfer information from the plant floor to supervisory and business computer systems. For example, most industrial plants are now using networked process historian servers and expert systems servers to allow business users to access real-time data from the DCS and PLC systems. There are also many other possible business/process interfaces, such as using remote X-Windows sessions from the DCS, or direct file transfer from PLCs to users' spreadsheets. Regardless of the method, each involves a network connection between the process and the business systems.

At the same time, there has been an explosion in the use of Ethernet and TCP/IP in industry for both process control and business networks. Most distributed control systems (DCS) now use Ethernet networking as a critical component of their system architecture, rather than the traditional proprietary industrial networks such as Data Highway or Modbus. Thus networks are increasingly Ethernet-based for both business and process systems.

The issue is that problems on the business network can be passed on to the process network through the business/process interface, and this can seriously impact production. Protecting the process system from external network problems is the focus of this paper.

II. EXAMPLES OF NETWORK PROBLEMS IMPACTING PRODUCTION

Many network designers divide the problems that can befall a process network into two general categories: accidental and deliberate. Accidental problems are typically caused by cabling and configuration errors or by user inexperience. Deliberate problems are those caused by individuals with malicious intent, such as disgruntled employees or network hackers. It has been our experience that accidental errors far out number the deliberate errors experienced in industrial environments but both should be addressed. Below we will look at a few examples of each type of error and how these errors have impacted process operations in North America.

Noise or Bad Packets: The most common network problem is the propagation of noise or bad packets through a plant network. For example, in February of 1996, a West Coast pulp and paper mill lost use of its entire business network as a result of a faulty network card in a workstation. Due to grounding problems, the network card started generating 1000 runt packets per second on the network (runts are packets that are so short they violate Ethernet rules). The network repeaters simply transmitted the packets to every section of the mill network, flooding the network and preventing any network activity. Fortunately mill production was not affected, due to some limited network protection already in place.

IP Address Duplication: As noted earlier, TCP/IP has become the most popular network protocol for industrial networks in the past five years. One of the requirements of TCP/IP is that every network device must have a unique IP network address. This address can either be manually entered into a computer's configuration or a central Dynamic Host Configuration Protocol (DHCP) server can automatically assign it. Either way, this number must be unique or problems will occur.

An example of this problem occurred in July of 1996 at the same paper mill as the previous example. Approximately one year prior, the mill had upgraded the profile controller on the #1 Paper Machine. This system used Ethernet and TCP/IP to communicate between the scanners and the main controller. It was also connected to the main mill network through a bridge so that profile information could be accessed by business applications. Some time after the installation, a network printer in another area of the mill was accidentally given the same IP address as the controller. Initially this did not cause difficulties, but shortly after a routine maintenance shutdown, the scanners started directing their data to the printer rather than to the controller. As a result, the paper machine could not be started for over six hours.

Broadcast Storms: Broadcast packets are messages that are directed to all the computers on a network rather than to a specific device. They may be generated by network servers advertising their services or by computers trying to locate other devices on the network. They are an important part of a properly functioning network and, in small quantities, have no negative impact.

In large quantities (what is referred to as a *Broadcast Storm*) broadcast packets can stop normal network operations. Each packet is perfectly valid on an individual basis, but demands that all network devices devote some CPU resources to interpreting it. Many common computers simply become overwhelmed if they receive too many broadcast packets in a short time span [1].

Two years ago a Saskatchewan industrial facility lost communications to the operator consoles on a steam plant DCS. The problem was believed to be caused by an incorrectly configured Windows 95 workstation in another mill area that generated high levels of broadcast packets. The DCS had to be removed from the mill network and remains disconnected to this day, preventing process data from being transferred to the business systems.

Deliberate Intrusion: Fortunately, deliberate intrusion of process control networks has been rare to date. However, as more mills attach to either the Internet or the corporate wide area network (WAN), the chances of being *hacked* are growing. Typically a hacker will attach to the mill network and attempt to locate possible

host computers he or she can invade. UNIX or VMS hosts (such as those used in many DCS systems) are popular because they have well known security holes that a hacker can exploit. Experienced hackers will use an automated security-scanning tool such as Security Analysis Tool for Auditing Networks (SATAN) software to check out an entire company network [2].

It is worth noting that system passwords only provide limited protection against hacking because most process control groups use very easy to remember (and easy to guess) passwords on their DCS or PLCs. At a recent ISA conference, the author was able to determine the passwords for the control system on a 16 site power generation system for a major Mid-Western US utility in less than 15 minutes.

Often the hacker is not an outsider with malicious intent but an employee doing something he or she shouldn't. A good example of this type of problem occurred this spring in a large East Coast paper mill [3]. The mill had just completed an upgrade of its paper machine, during which a number of engineers had been brought in from head office to assist with DCS commissioning. Everyone on the DCS commissioning team knew the passwords for the control system computers and when the project was completed, no one bothered to change them.

Trouble started about a month later when one of the head-office engineers decided he needed a good data source for an expert-systems experiment he was running. Using the company's wide area network (WAN), he was able to dial into the mill network from the corporate headquarters several hundred miles away. Once into the mill's business LAN, he was able to connect to the DCS through a link originally set-up to allow mill supervisors to view operators screens from their offices. He then loaded a small program onto one of the DCS graphics stations (a UNIX machine). This program asked all DCS devices to dump their data back to him once every five minutes.

All this would have worked fine, except that the engineer's new task would occasionally overload one of the DCS to PLC communications gateways, and it would stop reading the PLC data. This, of course, caused the machine operators great panic as they lost control of the motors controlled by the PLCs. Soon the electrical department was busy troubleshooting the PLCs. Meanwhile the head-office engineer had left the company to work for a competitor.

Eventually the problem was solved by an eagle-eyed mill engineer who noticed that the problems always occurred at intervals that were at multiple of five minutes. Suspecting that it might be software induced, he started to inspect all the tasks running on the DCS computers and found the offending task. Of course, by then the lost production in the mill had been substantial.

III. Managing Network Connections

From these examples it is clear that uncontrolled connections between the business network and the process network are undesirable. However, it is impossible in any modern facility to completely forbid the interconnection of process and business systems. Process data is too valuable to lock away on the plant floor. Instead, the process engineer must design a controlled connection point between business and process that will allow data to move through but prevent problems from being passed through. This connection point is known as a firewall.

Another way to look at this is to ask *"how do we deal with computers and devices that need to access both the business and process networks simultaneously?"* For example, many modern DCS configuration stations need to access the DCS but also need connections to CADD servers in the engineering office. If these stations have a single direct physical link to both the business and process networks, then this also implies an undesirable physical link between the all the process and business systems.

A commonly proposed solution to this dilemma is to keep the two networks separate by installing two network interface cards (NICs) on any workstation that needs access to both network systems. Unfortunately, this simple solution is not recommended.

A Windows NT workstation certainly can be configured so that traffic on one network interface is not passed to the other. Unfortunately, a single mouse click can enable the workstation to act as an uncontrolled router, so that it passes data between networks. The management effort to prevent this accidental routing can be significant, especially if the workstations are scattered around the mill.

A better solution is to select a single device to act a controlled access point between the two networks and then install the workstations on only one of the two networks. Any accesses to the other network would pass through this firewall and can be filtered for source, destination and task. Selecting and configuring this firewall is the key to effective security.

III. Communications Protocols

To understand the types of devices that might be used to protect a process control network, it is necessary to understand a little bit about communications protocols. Protocols are simply sets of rules that define how two machines communicate with each other. In a typical network, there will be dozens of protocols required simultaneously, each providing the rules for different communication functions such as flow control, error checking, message routing or even simple electrical signal to data conversion.

To help organize all of these protocols and understand how one protocol interacts with another, protocols are usually arranged into a layered model. Each layer in a model will group together protocols with related tasks. This way we can say that a specific layer has a specific function in a communications network. In addition, we say that each layer in the model provides a service to the layers above it.

The dominant layered model for organizing communications protocols is the one developed by the International Organization for Standardization (ISO). This is a seven-layer protocol model known as the Open Systems Interconnect Reference Model (OSI/RM). Figure 1 shows the organization of the seven layers in the OSI model and a few examples of where some well known protocols fit in. For our purposes of process network security, we can focus on understanding the bottom three layers [4]:

Physical - provides the standards for transmitting raw electrical signals over the communications channel. Physical protocols deal with the transmission physics such as modulation and transmission rates.

Data Link - the rules for interpreting electrical signals as data, error checking, physical addressing and media access control (which station can talk at any given time on the network).

Network - describes the rules for routing of messages through a complex network. Defines how to deal with network issues such as faulty lines and congestion.

The layers 4 through 7 are also important to a functional network but we don't need to deal with them at this time.

IV. Network Connection Hardware

Over the years four types of network devices have been defined to connect networks together:

Repeaters Bridges Routers Gateways

Each of these devices is designed to provide a connection between network protocols at a specific layer. In addition, as a result of their protocol conversion features, each of these devices can also provide some level of isolation between two networks with the same protocols. It is this feature that makes them important for network protection.

A. Repeaters

The repeater is designed to work at the physical layer, extending the length of a network by connecting two or more LAN segments and allowing conversion between cable types. They are mainly designed for regenerating electrical signals without filtering and provide little network protection. Most hubs and concentrators are repeaters.

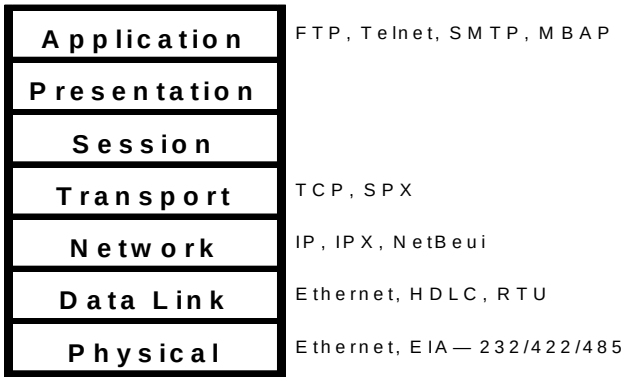

Figure 1: The OSI Reference Model for Communications Protocols. The seven OSI layers are illustrated along with some typical protocols assigned to their layers. Many well-known communications standards, such as Ethernet, span several layers.

The traditional repeater was a very 'dumb' device and could provide little or no isolation features. Today many Ethernet hubs are 'smart hubs' that monitor the traffic going through them and can cut off segments generating excessive errors, giving some limited network protection. For example, some smart hubs may have been able to localize the runt packets noted earlier to one network segment. However, as a general network protection device, the repeater is very limited.

B. Bridges

The function of a bridge is to connect separate networks together. These networks can be the same type (such as both Ethernet) or two different types (such as Ethernet and Token-Ring).

Bridges work at the data link layer of the OSI model, recording the physical addresses of the nodes on each network connected to the bridge and then allowing only the necessary traffic to pass through the bridge. When a message is received by the bridge, the bridge reads the packet and determines the destination and source addresses of the message. If the source and destination networks are different, the packet is passed through. If both addresses are from the same network, the message is not passed on.

This feature of a bridge is very important for controlling network loading. If the traffic between two computers is over-loading a network, a bridge will prevent the traffic from propagating onto other networks and over-loading those networks as well. In addition a bridge will check the physical integrity of each message and block bad messages from crossing. For example, a bridge isolating the administration network from the process control network would prevent heavy email traffic between accounting stations from tying up the process network. It would also stop noise-corrupted packets from a computer in engineering from getting into the process network. A bridge would certainly have prevented the runt packets in the first example from spreading through out the mill.

When a mill network is sub-divided into separate networks joined by bridges, we say the mill network is divided into separate *collision domains*.

C. Routers

Routers operate at the network layer of the OSI protocol model and work with packets based on network protocols they contain. They forward messages through complex networks (such as the Internet), selecting the best possible route based on criteria such as availability, loading, cost and speed.

Routers are intelligent devices used to divide networks logically rather than physically. For example, an IP router can divide a network into various subnets so that only traffic destined for particular IP addresses can pass between segments. Limiting of broadcasts packets to small *broadcast-domains* is a common use for routers.

Another router feature is filtering. To make intelligent routing decisions, a router needs to know a lot about the message it is handling. For example, live video over a network would need a fast route but email could go over a slow but inexpensive route. As a result, routing protocols such as IP contain information about type of packet (e.g. email, file transfer, telnet, video, etc.) and its ultimate source and destination.

This feature can be very useful for security because it allows us to use the router to filter out certain types of messages from ever entering a control network. A router connecting a process network to a business network might be configured to filter out all messages entering the process network except X-windows traffic. All email, file transfers or telnet sessions from the business side could not enter the process network. This filtering of network traffic through a router is known as *firewalling* a network.

D. Gateways

Gateways provide full seven-layer protocol support. They are used to connect to completely differing systems (such as Provox DH II to DEC or Novell to IBM). They can also be used to provide application layer conversions, say between two different email systems.

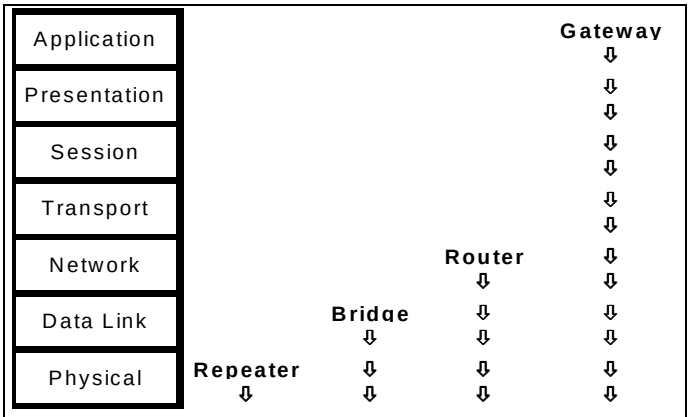

Figure 2: The Major Network Devices As They Fit Into The OSI Model - The repeater, bridge, router and gateway are defined by the protocol layers they work with. Each can interpret their key layer plus any layer underneath.

E. Switches - A New Device

A new network device that has attracted a lot of attention recently is the network switch. A switch is basically multi-port bridge (a layer-2 switch) or router (a layer-3 switch) with a very high-speed backplane. Each port connects to an independent network that operates as its own collision domain or broadcast domain. The high-speed backplane transfers the inter-port messages between ports.

V. NETWORK DESIGN OPTIONS

A. Traditional Network Architecture - Bridging

Traditionally, local area network designs were based on flat networks connected by simple bridges to a backbone of a similar protocol [5]. For example, Thicknet Ethernet (10BASE-5) often acted as a backbone running between major centres in an industrial facility. Bridges were attached to this backbone and these allowed the connection of department networks that were also based on Ethernet. The bridges provided isolation of the network traffic and some limited protection against the propagation of bad packets throughout the network. Figure 3 shows a typical bridge-based architecture.

This design was especially true for process control networks, where, until five years ago, the practice was to run completely separate systems for the mill control network and the MIS network and connect them together with a bridge. The problem with this technique is that it offers no protection from broadcast packets and no security features.

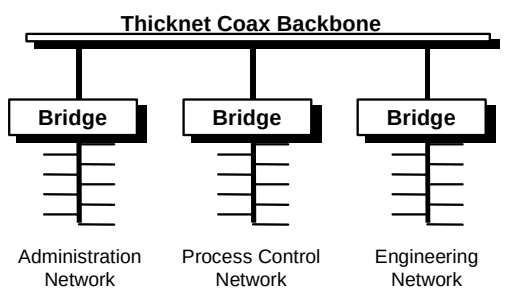

Figure 3: Traditional Bridged Network Design. With this architecture, department networks are connected by simple bridges to a network backbone of a similar protocol. The bridges provide some limited protection against the propagation of bad packets and heavy traffic but offer no broadcast storm control.

B. Current Network Architecture - Routing Switches

The current network design strategy is to combine switches and routers in either a single device called a routing or layer-three switch. At a minimum, these devices offer basic packet security and containment of broadcast storms in a very high throughput device. More typically, layer-3 switches have the ability to define filters based on address, IP subnet, protocol or application.

The use of routing switches is recommended for industrial networks because it combines the broadcast storm containment and basic security of routers with the speed of switches. Conventional routers are too complex and costly for most LAN applications. They are better suited for wide area networks (WAN) environment where their increased cost and complexity is justified by the expense of WAN bandwidth and security issues of external access. Conversely, bridges are fast enough but do not offer sufficient security against broadcast storms or network intrusion.

If a traditional IP router is used to separate network areas we say that the network is divided into subnetworks. If a Layer-3 switch is used we typically speak of dividing the network into a number of Virtual Local Area Networks (VLANs). In both cases, the router or layer-3 switch is the center point for all the network traffic. When two devices are defined as being on the same subnet or VLAN, the switch passes through messages with no filtering, just as if the devices were on the same physical segment. However, if two devices are not on the same VLAN, then the switch runs the message through its filtering software, passing or blocking the message as appropriate [6].

It is important to remember that the router or switch can only filter traffic that passes through it. It cannot separate two devices if they are physically wired to the same segment. Thus if it is important to filter traffic between two different groups of devices, make certain that they are attached to different switch or router ports.

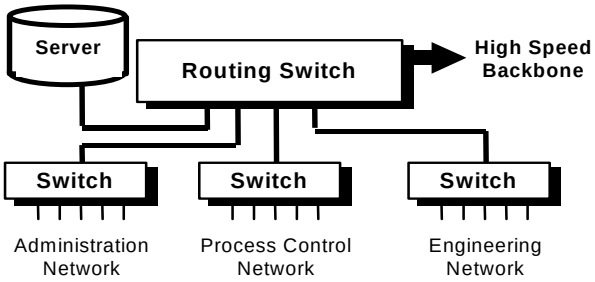

Figure 4: Network Design Using Routing Switches. In this design all inter-department network traffic goes through one or more routing switches. This allows network administrators a single point to manage network protection and security using VLANs and filters.

VI. IMPLEMENTING SECURITY POLICY

Once a router or Layer-3 switch is installed between business and process networks, it is necessary to implement a security policy. This will be a set of rules that dictate who is permitted to access various areas of a network and what type of action is permitted once this access has been granted. Translated into networking terms, these are a set of traffic filters that define the network addresses, protocols and port numbers that data packets must conform to before they are passed into a secure area. These traffic filters may monitor one or all of the subnet connections and may be applied to either incoming or outgoing traffic The following are types of traffic filtering that are typical available for an IP network.

- Filtering on specific IP addresses.
- Filtering on IP address ranges. This utilizes IP subnet masks to identify a specific address range.
- Filtering TCP/UDP port numbers. This allows filters to be application specific. For example Telnet traffic always uses the TCP incoming port number of 23.
- Matching can be performed on packet source addresses or destination addresses.

All filtering may be of the form “Permit” or “Deny”. Depending on the desired security strategy, it may be best to permit access to a few specific types of packets and deny everything else.

These rules may be applied in a number of ways. Using addresses and masks it is possible to set up the following combinations:

- Many addresses to many addresses.
- One address to many addresses.

- Many addresses to one address.
- One address to one address.

It is also possible to filter on the application of packet that is being passed to the process networks. For example, Telnet and HTTP may be blocked but FTP might be permitted. To achieve this, the engineer will need to determine the TCP/UDP port numbers of all applications that should be passed through and then block all others.

VI. A Case History – DCS Configuration Stations on the Business Network

A major pulp mill in Quebec was up grading the DCS system to use NT-based workstations for both operator consoles and DCS engineering configuration stations. The operator consoles presented little security risk as they were located on the process network and could only run the DCS graphics application.

The engineering configuration stations were much more of a problem. First of all, they needed to communicate with the CADD and maintenance servers on the business network as well as the DCS. In addition, the process engineers wanted the DCS configuration stations to be located in the engineering area, far from the process network. After discussions with the DCS vendor, it was decided that they should be attached directly to the business network rather than the process network. This would require a network security policy that would the central routing switch to block other traffic to the DCS servers but allow configuration traffic from these machines.

To start off, IP addresses were carefully assigned to make the filtering as easy as possible. The DCS controllers, operator consoles and servers were assigned an IP address range of 10.4.1.1 to 10.4.1.254 with a subnet mask of /24. The DCS servers were given an address of 10.4.1.1 and 10.4.1.2. The Engineering Workstations on the business network had an IP address range of 10.5.1.1 to 10.5.1.254 with a subnet mask of /24. Protocol analysis showed that the TCP port numbers were 6000 connecting to the DCS and between 1152 and 1160 in response (See Table 1).

Action	Dst Port #	Src Port #
Log-in to DCS server	6000	1153, 1154, 1155, 1156
Perform DCS diagnostics	6000	1154, 1156, 1157
View Graphics	6000	1154, 1156, 1159
Print Report	6000	1154, 1156

Table 1: SampleTCP Port numbers recorded during configuration sessions to the DCS server

For most routing hardware, the filtering rules are generally of the form:

[Permit/Deny] [Source Host/Mask] [Destination Host/ Mask] [TCP Socket = ???]

Therefore the rules that were applied were:

[Permit] [10.5.1.0/24] [10.4.1.0/31] [TCP Socket = 6000] *Permit engineering workstation traffic to enter the DCS network if directed to the two servers and using TCP port # 6000.*

[Permit] [10.4.1.1/32] [10.5.1.0/24] [TCP Socket >= 1152 & TCP Socket <= 1160] *Permit only the DCS Server #1 traffic to enter the business network. if directed to the engineering workstations and using TCP port # between 1152 and 1160 .*

[Permit] [10.4.1.2/32] [10.5.1.0/24] [TCP Socket >= 1152 & TCP Socket <= 1160] *Permit only the DCS Server #2 traffic to enter the business network. if directed to the engineering workstations and using TCP port # between 1152 and 1160 .*

[Deny] [0.0.0.0/0] [10.4.1.0/24]
This final rule denies all other traffic into the DCS network.

This example illustrates three types of filtering. The first is filtering on a range of addresses using a subnet mask. The second is filtering on a host using the complete address. The third to the filtering of each packet for TCP socket number.

Each packet entering or leaving the DCS network through the switch would be checked against each of these rules in turn. As soon as a rule is satisfied the packet is either passed through or deleted, depending on the rule. This is why the final DENY statement is required.

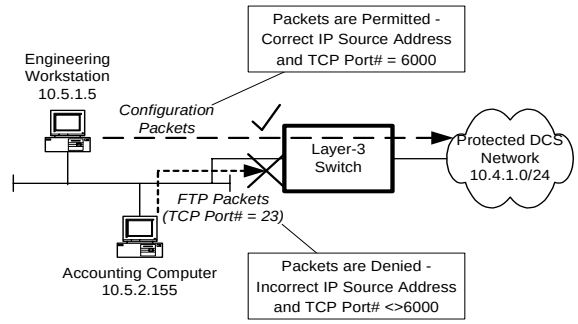

Figure 5: Packet Filtering Using Routing Switches. All inter-department network traffic goes through one routing switch that allows network administrators a single point to manage network protection using IP security filters. In this case only messages originating at computers with the 10.5.1.0 subnet and using TCP socket 6000 would be allowed into the DCS network

VII. Conclusions

The experience at this and other mills clearly show that a network security design is important in any plant integrating the process and business systems. Conventional bridge-based designs or dual carded workstations that are used without a strategy won't prevent many serious network problems. The stages in a good design include installation of a single business/process interface point (usually a Layer-3 switch), careful assignment of IP addresses and development of appropriate traffic filters.

References

[1] Cisco Systems Inc., Internetwork Design Guide - Appendix E, 1997

[2] Semeria, C., *Internet Firewalls and Security*, 3COM Technical Paper, pp. 4-5, 1996

[3] Byres, E. J., *Network Secures Process Control*, InTech, Instrument Society of America, pp. 92-93, October 1998

[4] Tanenbaum, A., *Computer Networks*, 2nd Ed., Englewood Cliffs, NJ: Prentice-Hall, 1989.

[5] Gohn, B. and Howe, G., *High Function Switches* 3COM Technical Paper, pp. 2-3, 1996

[6] Mandeville, R. and Newman, D., *VLANs: Real Virtues*, Data Communications, May 1997, P. 82

About the Author:

Eric J. Byres is a Professional Engineer and research manager of the Internet Engineering Lab at the British Columbia Institute of Technology, one of North America's leading research facilities in the field of industrial cyber-security. He is also a faculty member in BCIT's School of Electrical and Electronic Technology. For the past 14 years, he has specialized in data communications and controls systems in industrial environments, focusing on industrial Ethernet research and network security design. He also has extensive experience in DCS and PLC system design and commissioning.

Eric was a founding member of the Instrument Automation and Systems Society (ISA) TECH/EXPO Networking and Fieldbuses steering committee and the Chair of the ISA Industrial Ethernet Conference in 2000 and 2001. He is a board member of the Industrial Automation Open Networking Association and is on the NIST committee developing industrial network security standards. He holds the Advanced Systems Institute (ASI) fellowship for industrial network security research until 2005. In 1999, he was the winner of the Best Paper Award at the IEEE Pulp and Paper Industrial Applications Conference and in September 2000 he won the IEEE Outstanding Industry Applications Article award for his paper on network security.