Eric J. Byres

DESIGNING SECURE NETWORKS for Process Control

In most industrial plants, there is a strong drive to provide business applications with some access to real-time data generated from process control systems. Often, this has been in the form of a process historian database server connected to both the distributed control systems/programmable logic controller (DCS/PLC) systems and the business users. It can also take many other forms such as remote X-Windows sessions from the DCS, or direct file transfers from PLCs to users' spreadsheets. Regardless of the method, it involves a network connection to both the process and the business side.

These network connections are increasingly Ethernet based on both sides, rather than proprietary industrial protocols. At the same time, most control systems now use Ethernet networking as a critical component of their system architecture. This can be for controller-to-controller communications, controller-to-operator console communications or even I/O-to-controller communications. Losing any of these links will directly impact production.

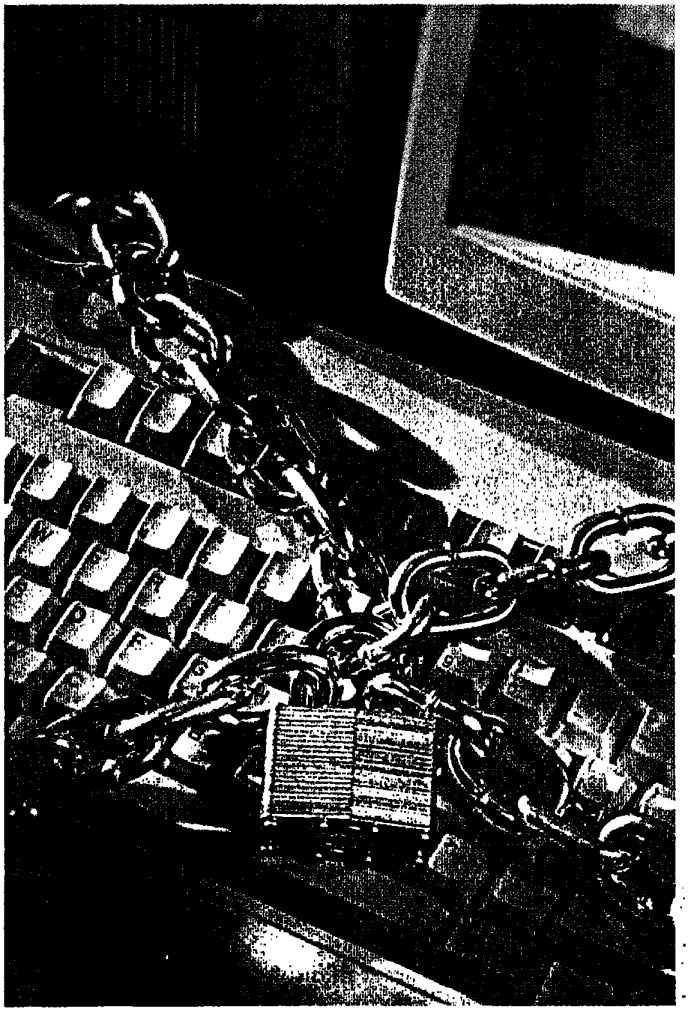

Eric J. Byres is with Artemis Industrial Network Design and Training. He is a Member of the IEEE. *In its original form, this article won the Best Conference Paper Award at the 1999 IEEE Pulp & Paper Industry Applications Conference.*

1077-2618/00/$10.00©2000 IEEE

The traditional repeater was a very "dumb" device and could provide little or no isolation features.

The issue is that problems on the business network can be passed on to the process network through this process data link, seriously impacting production. Protecting that process system from external network problems is the focus of this article.

Network Problems that Impact Production

Problems that can attack a process network from the outside world can be divided into two general categories: accidental and deliberate. Accidental problems are typically caused by cabling and configuration errors or by user inexperience. Deliberate problems are caused by individuals with malicious intent, such as disgruntled employees or network hackers. It has been our experience that accidental errors far outnumber the deliberate errors experienced in industrial environments, but both should be addressed. We will explore 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 mill network. For example, in February 1996, a west coast pulp and paper mill lost the 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

TCP/IP has become the most popular network protocol for mill networks in the past three 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.

One example of a problem occurred in July 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].

In 1998, a Saskatchewan industrial facility lost communications to the operator consoles on a steam plant DCS. The problem was believed to have been 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, the 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 to 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's 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 sys-

tem for a major mid-western U.S. utility in less than five 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 at 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 the 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.

The trouble started about a month later when one of the head-office engineers decided that he needed a good data source for an expert-systems experiment he was running. Using the company's 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 the 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 bothering to get the PLC data. This, of course, caused the machine operators great panic as they lost control of the motors controlled by the PLCs. Soon afterward, 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.

Communications Protocols

To understand the methods available 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 or data highway, 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.

Often the hacker is not an outsider with malicious intent, but an employee doing something he or she shouldn't.

To help organize all of these protocols and understand how one protocol interacts with another, they are usually arranged in a layered model. Each layer groups 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). Fig. 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 the purpose of process network security, we will 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—has 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 messages through a complex network. Defines how to deal with network issues such as faulty lines and congestion.

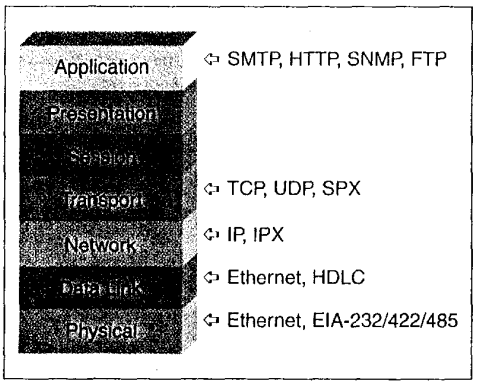

Fig. 1. The OSI reference model for communications protocols. The seven OSI layers are illustrated along with some typical protocols assigned to their layers. Many well-known communications standards, such as Ethernet, span several layers.

A switch is basically a multiport bridge (a layer-two switch) or router (a layer-three switch) with a very-high-speed backplane.

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

Network Hardware

It is important to understand that the OSI layer specifications are functional only; what to do is defined, but how to do it is not. Thus, two protocol families that are "ISO compatible" won't necessarily communicate directly. For example, in the 1980s, IBM, DEC, and the Internet each had very different methods of routing messages through a network, but all three might have said that their routing protocol would fit into the network layer of the OSI model.

Since there were many different possible protocols at any layer in the OSI model, some means of connecting networks using different protocols was needed. Four classes of network devices are defined to address this issue:

- Repeaters,
- Bridges,
- Routers, and
- Gateways.

Each of these devices is designed to provide a connection between different 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. We will look at each device to show how it might be used to protect a network.

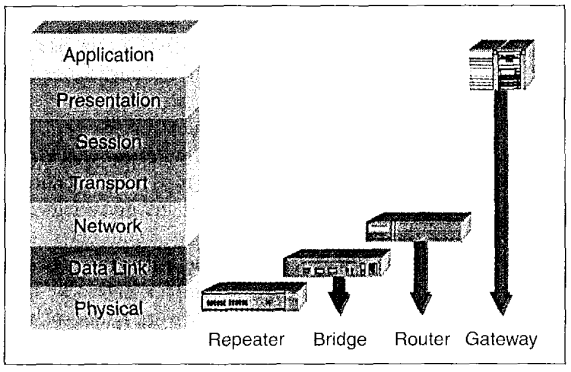

Fig. 2. The major network devices as they fit into the OSI model. The repeater, bridge, router, and gateway are defined by the protocol layers with which they work. Each can interpret their key layer plus any layer underneath.

Repeaters

The repeater is designed to convert between two physical layer protocols, provided that the upper layer protocols are the same. For example, a repeater could convert between twisted-pair Ethernet and fiber Ethernet, but not between twisted-pair token-ring and fiber Ethernet. It can also be used to connect two or more LAN segments to extend the length of the network by repeating the signal. Most hubs and concentrators are repeaters.

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.

Bridges

The function of a bridge is to connect separate networks. 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

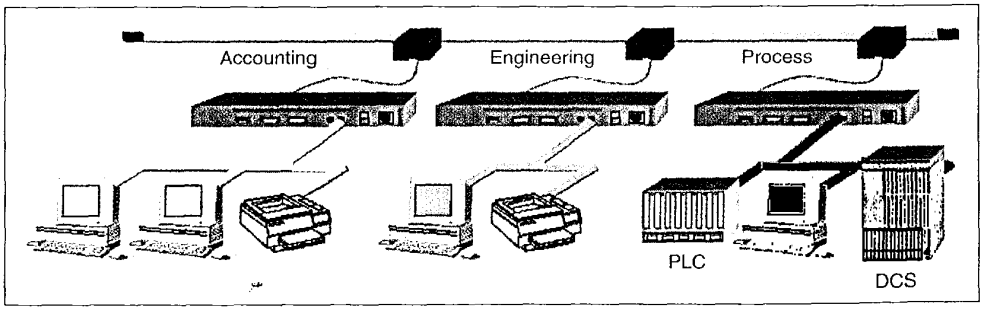

Fig. 3. Network design using routing-enabled switches. In this design, all interdepartmental network traffic goes through one or more routing-enabled switches. This allows network administrators a single point to manage network protection and security using VLANs and filters.

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 overloading a network, a bridge will prevent the traffic from propagating onto other networks and overloading them 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 in accounting 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 from spreading through out the mill in the previous example. When a mill network is subdivided into separate networks joined by bridges, we say the mill network is divided into separate *collision domains.*

Routers

Routers operate at the network layer of the OSI protocol model and work with packets based on the network protocols they contain. More importantly, routers help 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 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 the 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 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.

Gateways

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

Switches: New Devices

A new network device that has attracted a lot of attention recently is the switch. A switch is basically a multiport bridge (a layer-two switch) or router (a layer-three 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 interport messages between ports.

Network Design

Traditional Network Architecture: Bridging

Traditionally, LAN designs were based on flat networks connected by simple bridges to a backbone of a similar protocol [5]. For example,

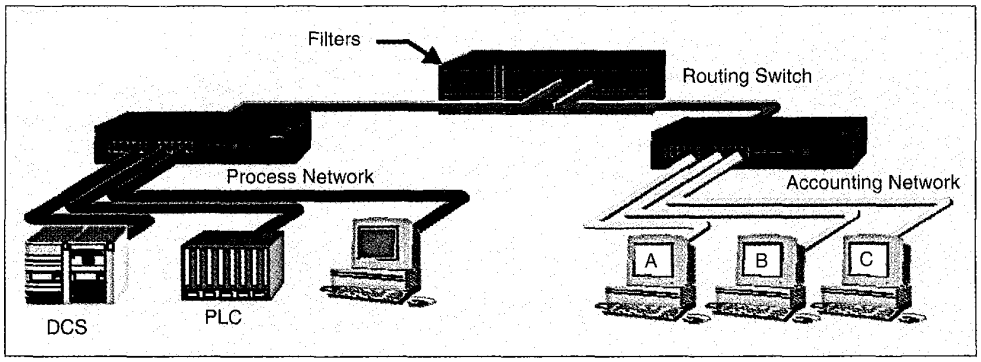

Fig. 4. Traditional bridged-network design. With this architecture, department networks are connected by simple bridges to a network backbone of a similar protocol. The bridges provide some limited protection against the propagation of bad packets and heavy traffic, but offer no broadcast storm control.

Thicknet Ethernet (10 Base-5) often acted as a backbone running between major centers 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. Fig. 4 shows a typical bridge-based architecture.

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

Network Architecture: Routing-Enabled Switches

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

The use of routing-enabled 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 WAN environments 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.

Virtual Local Area Networks

When two devices are attached to the same network segment, every message sent by one of the devices will be directly received by the other device, without any filtering. On the other hand, if two devices are connected via a router, then every message will be subject to some degree of filtering. The first case is obviously faster since there is no intermediate processing involved. It is ideal when the devices are designed to work with each other, such as DCS components. The second case, however, offers more security.

In an ideal network, we would be able to wire each group of "matched" devices on their own private network, and connect all the groups through routers or routing switches. In the real world, however, devices that ideally should be on the same network segment may be scattered over a wide area, making wiring them together very expensive. Even more serious, two devices may each be designed to talk directly to a third, but not to each other. For example, the two devices may be a process workstation and a business workstation, while the common third device is a process information server. Clearly, a direct physical link between each of the workstations and the server also implies an undesirable physical link between the process and business workstations.

The solution to this dilemma is the virtual LAN (VLAN). In the VLAN, a routing switch is the center point for all the network traffic. When two devices are defined as being on the same 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 switch can only filter traffic that passes through it. It cannot separate two devices if they are physically wired to

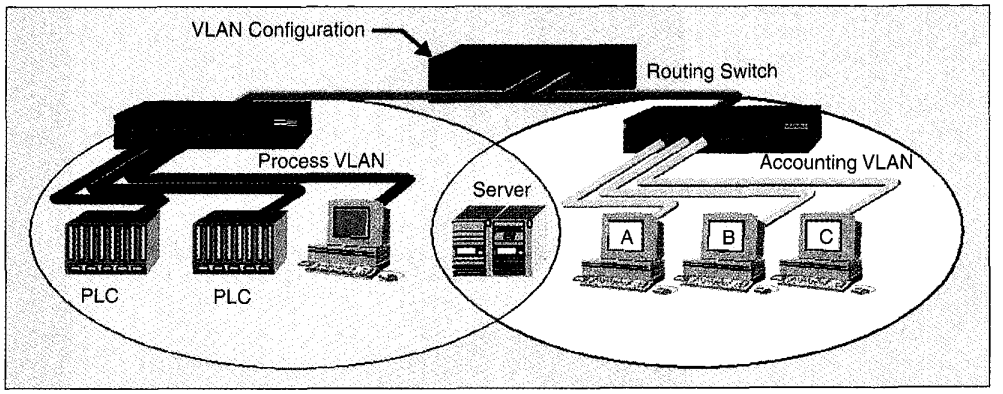

Fig. 5. Two separate VLANs both containing the same process information server. In this example, separate VLANs have been set up for both the business users and the process control users. Both VLANs contain the process information server, allowing both groups to access the server, yet forming a secure separation between process and business networks.

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 ports.

A Case History: The Tembec Mill-Wide Information System

The Tembec pulp mill in Temiscamming, Quebec, intended to install a mill-wide process information system to give up-to-the-minute process information to management both inside and outside the mill. Before any system could be installed, however, a networking strategy was needed that would provide a reliable and secure connection between the process control systems and the mill-wide business network.

The network at the Tembec mill consisted of a site-wide fiber optic ring connecting seven Ethernet switches. The switches provided approximately 40 Ethernet segments connecting approximately 250 process and business users. Existing DCS and PLC systems already utilized some Ethernet networking, but were isolated from the rest of the mill.

The design of the network protection centered on upgrading the existing bridging switches to routing-enabled switches and then the development of VLANs and filters. After discussions between the mill staff and various equipment suppliers, it was decided that the following design strategy should be used:

- One general business VLAN and two process VLANs would be created initially. One VLAN would encompass the DCS and process data collection and one would be for PLC programming. A separate MMI network for VLAN work may be created in the future.
- The first level of VLANs configured should be IP routing-based VLANs.
- The second layer of network security would be to install protocol filters to restrict traffic between the business and process networks to IP traffic only. All DECnet and Netbeui traffic should be prevented from crossing between VLANs.

It was also decided that firewall protection from possible Internet hackers be done by the IS group at the mill to the Internet connection point. Internal security would be done through application-based VLANs, but these would not be installed until the process information system was more stabilized and the mill networking staff more familiar with configuring the switches.

This initial configuration was completed in December 1997 and was tested at that time. Interestingly, while working on the VLAN configuration, IP broadcast storms were noted throughout the network. Typical traffic captures showed approximately 3100 broadcast packets per minute and 600 multicast packets per minute. This IP broadcast traffic often accounted for over 80% of the network load in many areas.

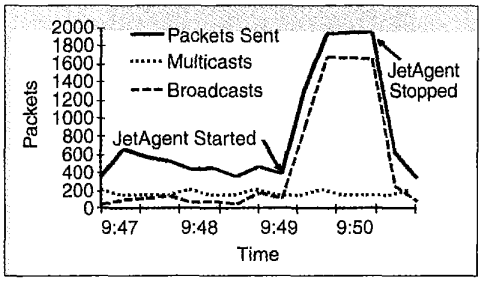

Fig. 6. Network traffic graph showing a broadcast storm on an unprotected segment. The printer management software is started and stopped to show that it is creating most of the segment traffic. There probably was sufficient broadcast traffic to impact the DCS if it had been connected at the time.

Further analysis determined that most of the broadcast packets were coming from the printer management computer. A printer management package was determined to be the problem and was disabled. Fig. 6 shows the effect on network traffic as this application was started and stopped.

The number of broadcast packets would have seriously impacted the DCS operator network if it had been connected to the mill network by a hub or bridge. Even an incorrectly configured router would not have helped. However, with the VLANs in place, no broadcast packets entered the process network.

Conclusions

The experience at Tembec clearly showed that a network security design is important in any mill that integrates their process and business systems. Conventional bridge-based designs or routers that are used without a strategy won't prevent many serious network problems. Tembec also showed that the design is a phased process that needs to be implemented in stages. There is no single solution that will address all network security issues faced by mill process control networks.

References

[1] *Internetwork Design Guide*, Appendix E, Cisco Systems Inc., San Jose, CA, pp. E1-E2, 1997.

[2] C. Semeria, "Internet firewalls and security," 3COM Technical Paper, Santa Clara, CA, pp. 4-5, 1996.

[3] E.J. Byres, "Network secures process control," *InTech*, pp. 92-93, Oct. 1998.

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

[5] B. Gohn and G. Howe, "High function switches," 3COM Technical Paper, Santa Clara, CA, pp. 2-3, 1996.

[6] R. Mandeville and D. Newman, "VLANs: real virtues," *Data Communications*, vol. 12, no. 6, p. 82, May 1997.