Implementing Firewalls for Modern Substation Cybersecurity

Dwight Anderson and Nathan Kipp
Schweitzer Engineering Laboratories, Inc.

Presented at the
12th Annual Western Power Delivery Automation Conference
Spokane, Washington
April 13–15, 2010

Implementing Firewalls for Modern Substation Cybersecurity

Dwight Anderson and Nathan Kipp, *Schweitzer Engineering Laboratories, Inc.*

***Abstract*—Firewalls are ubiquitous cybersecurity tools that are present in many computer operating systems and communications networks. This paper offers a high-level view of the basic operation of a firewall. It also shares how a firewall increases the security posture of a modern substation. Further, this paper identifies some considerations for firewall implementation and operation as a cybersecurity tool for modern power substation design that needs to support new regulatory requirements.**

I. Introduction

Power utility systems deliver information to a wide range of users in near real time and automate several tasks that streamline operations and performance. These systems communicate over many types of networks. Today, all over the world, power system automation experts are installing Ethernet local-area networks to increase the flexibility and capabilities of their systems and to reduce costs.

Historically, Ethernet operated with the idea that all devices on a network are trustworthy. That is, networked devices can be trusted with any information they receive from the network, and networked devices will not attempt to cause harm to the network or other networked devices. The idea that all Ethernet devices are trustworthy worked well when there was limited access to computer networks. These ideas promoted communications efficiency and contributed to the speedy development of newer communications standards.

As computer networks became more accessible to the general public, the security weaknesses built into Ethernet networks became apparent. We now know that not all devices or persons with access to a computer network should be trusted. This is especially true when we cannot positively identify who has access to our networks. The network firewall is one tool that mitigates the risks inherent in Ethernet communications.

One of the origins for the term "firewall" comes from the firewall that is in automobiles between the engine and passenger compartments. Its purpose is to confine a fire to the engine compartment, preventing the fire from progressing to the driver and passengers. An automobile firewall provides holes for functions, such as steering, throttle control, and braking. Similarly, a network firewall restricts unauthorized traffic from flowing on a network segment but allows authorized data to proceed. Network firewalls operate at network boundaries where communications from different networks meet (e.g., where a corporate intranet interfaces with the Internet).

To understand how network firewalls operate, we need a good understanding of Internet Protocol (IP) communications. "Warriors of the Net" is an instructive video that tracks an IP packet as it travels through the Internet [1].

Firewalls come in two types, network-based firewalls and host-based firewalls. This paper focuses specifically on network-based firewalls, but all the ideas presented here apply to host-based firewalls as well.

II. Types of Firewalls

Firewalls employ rules that permit, limit, or deny the transmission of packet-based communications and, in doing so, reduce the exposure of communications networks to external threats. The rules that permit traffic appear as holes in the firewall, similar to the holes in an automobile firewall, and allow authorized network traffic for these services to pass. Rules that deny access instruct the firewall to drop traffic not authorized to progress through the firewall. Also, some firewalls that drop data packets often create an alarm or log file that notifies the user and/or administrator of the actions. Be aware of and investigate chatter resulting from an excessive number of dropped packets. Fig. 1 shows a typical network diagram with firewall placement. The firewall protecting the supervisory control and data acquisition (SCADA) data traverses a wide-area network (WAN).

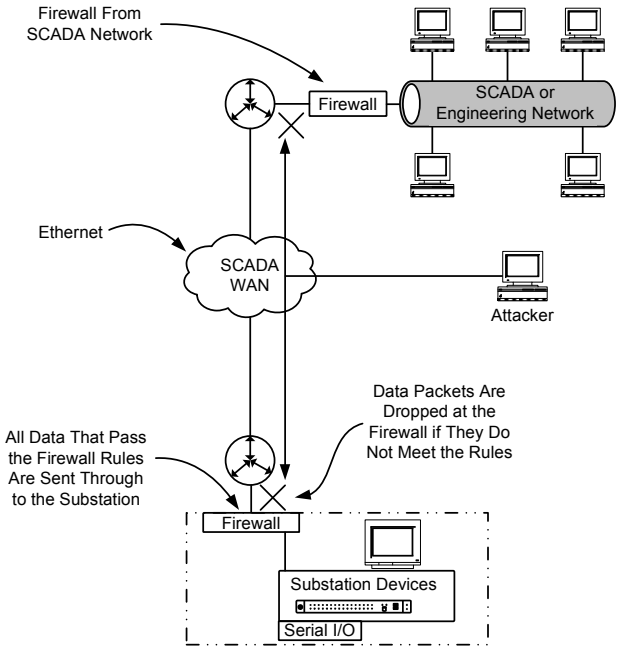

Fig. 1. Typical network topology with firewalls

As with any security tool, firewalls require an in-depth understanding of network design. Unintentionally or inaccurately changing a firewall rule can impede important network traffic or, conversely, permit or allow unwanted traffic to pass.

The following are common types of network-based firewalls, as shown in Fig. 2:

- Packet filtering
- Stateful inspection
- Application gateway
- Network proxy

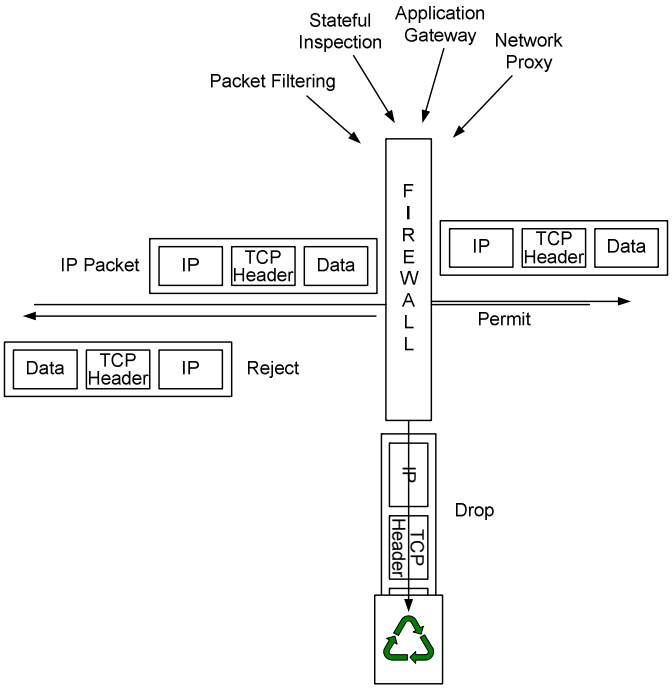

Fig. 2. Four types of firewalls

Packet-filtering firewalls examine the IP addresses and port numbers of traversing packets. The packet-filtering firewall compares these values against a set of programmed rules to determine the action to perform on each packet. Packet-filtering firewalls are fast but have limited ability to secure the data passing through the wall. They are susceptible to IP spoofing attacks.

Stateful inspection firewalls (sometimes referred to as session-based firewalls) have all the features found in packet-filtering firewalls. Additionally, they dynamically create rules based on the state of established connections. Once the connection is established, no further packets that are part of that connection need examination, making stateful firewalls very fast. Stateful firewalls are also good at protecting dynamic protocols, such as File Transfer Protocol (FTP).

Application gateways work at the application layer, Layer 7, of the Open System Interconnection (OSI) model. This layer holds the transmitted application data. Application gateways filter all traffic from rules that apply to the payload of the data packet. These firewalls provide the highest level of security but at the cost of speed.

Network proxies intercept all messages entering or leaving the network. Administrators often combine network proxies with other types of firewalls, such as application gateways. Network proxies are able to hide the true addresses of internal network devices.

It is important to consider the types of communications that traverse a network boundary when implementing a firewall solution. There are three core protocols, in addition to IP, used in Ethernet communications:

- Internet Control Message Protocol (ICMP)
- User Datagram Packet (UDP)
- Transmission Control Protocol (TCP)

Networked devices use ICMP when performing diagnostics and sharing connection status information. The ping and traceroute functions of ICMP provide network administrators with the tools to diagnose communications errors. Likewise, these functions give attackers the ability to map target networks. Firewalls should have the ability to easily enable and disable ICMP messages. Firewalls are important for network administrators to efficiently perform their work, but they should not be allowed to provide attackers a way to map the system.

UDP is a connectionless transport protocol that is similar to sending a post card. As the information or post card is sent, the protocol provides no confirmation that the message is received.

TCP is the primary transport protocol used in computer networks. It differs from UDP in that TCP is connection-oriented. Namely, all messages receive acknowledgments and form the basis for very reliable communications. This shapes the basis of the state of a connection.

TCP connections form a start state with a three-way handshake, as shown in Fig. 3. A TCP handshake involves an initiator sending a SYN packet to a responder. The responder then replies with an acknowledgment (ACK) and a SYN packet of its own. The initiator then acknowledges the responder's SYN packet to establish the TCP session. The state of a connection is a helpful resource for firewalls to use to protect the data.

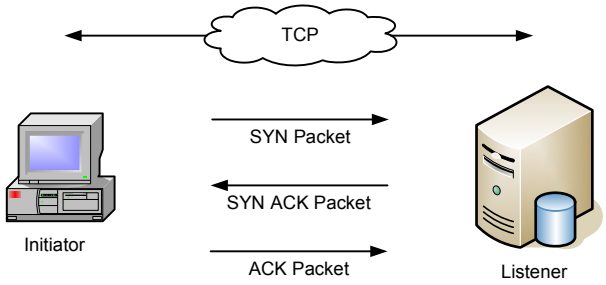

Fig. 3. TCP three-way handshake

III. Considerations of Firewalls for Substations

A modern substation should contain, at a minimum, a stateful, deny-by-default firewall to prevent unauthorized incoming (ingress) and outgoing (egress) communications. It is important to filter ingress communications to prevent attacks from external sources. Filtering egress communications is just as important to prevent the leakage of sensitive information.

A stateful, deny-by-default firewall examines the source and destination addresses and port numbers of all incoming or outgoing packets. The firewall compares this information to a set of user-configurable rules to determine what action to perform on that packet. There are three actions a firewall might take: forward, drop, and reject. A packet is forwarded when it is allowed to progress from one side of the firewall to the other. Rejected packets are not routed to their destination, and a message indicating the rejection is sent to the packet source. Packets that are dropped provide no response to the source and do not progress through the firewall. Dropping packets rather than rejecting packets is beneficial to hide the presence of the firewall and the fact that a protected network exists at all.

Firewall rules for modern substations should take the deny-by-default approach and only allow packets to progress through the firewall when holes are opened in the firewall by authorized administrators. This means that the firewall does not allow any traffic to pass until someone creates rules that specifically authorize certain types of traffic to pass. Explicitly identifying authorized traffic helps to ensure that the firewall denies all unwanted traffic by default. This mitigates the chance of misconfiguration that might allow unauthorized traffic.

Monitoring the state of a network connection enhances how a firewall mitigates risk. The state of a network connection is dependent on a number of variables arising from the situation or application. The transport protocol, network protocol, and application all affect connection states. There are several states for TCP connections; firewalls are most concerned with the following three:

- SYN sent: a SYN packet has been sent to establish a connection. A response has not yet been received. This state applies only to the initiator.
- SYN received: a SYN packet has been received from the initiator. The responding SYN packet has not yet been acknowledged. This state applies only to the responder.
- Established: a three-way handshake has been successfully completed, and the TCP session is active.

A stateful firewall monitors the states of connections and compares them to dynamic rules in order to determine if data should proceed. Stateful firewalls provide greater ability to protect against risks because they look at more packet information than just the source and destination IP addresses and port numbers, as shown in Fig. 4. A stateful firewall follows the state of a connection, adding greater security to the communications link.

Firewalls must have the ability for users to manage their own rule sets. Predefined manufacturer rules are useful and provide for quick setup, but alone do not give the granularity of control needed to secure critical infrastructure systems. Predefined manufacturer rules may open ports or services that have no merit in a control system operation.

So far, we have focused on what a firewall is and the basics of firewall operation. The following section focuses on specific rules that a modern substation firewall should implement in order to increase the overall security of the substation.

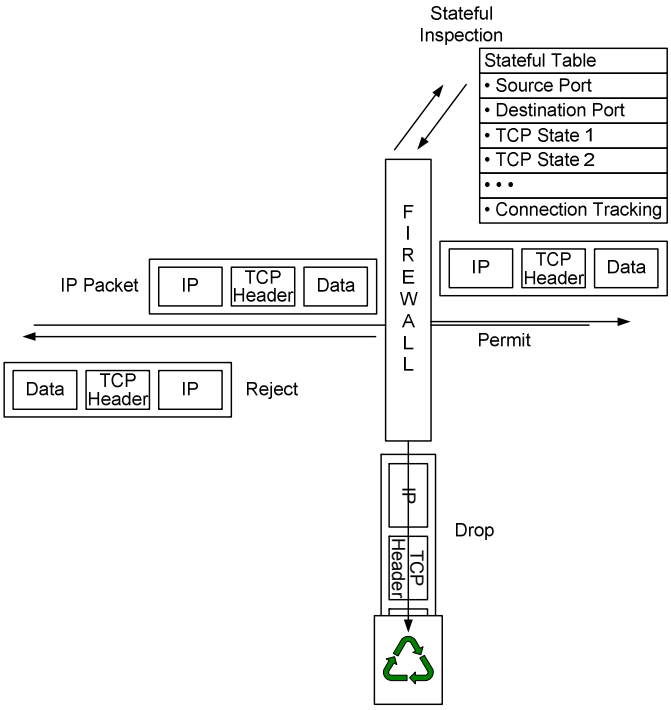

Fig. 4. Stateful inspection firewall

IV. Firewalls, Ports, and Services

There are 65,535 available networking ports on computers and network devices that reside in the modern substation. Many services are assigned to specific computers or network ports for ease of management and interoperability. Some of the network management protocols provide services, such as printer and file sharing or network device management. These services, such as *bootp/bootpc*, Simple Network Management Protocol (SNMP), and remote procedure call (RPC), lend themselves to compromise because they provide so much control over networked devices. This paper recommends that substation network designers block these types of network management services from operating across untrusted connections.

Some firewalls only allow cryptographic data to pass. They help to mitigate the risks associated with network management protocols. Firewalls that allow only cryptographic data to pass protect data from unauthorized entities. They also ensure that any data entering the network come from a trusted source.

Previously, we discussed ICMP messages. Modern substation firewalls should prevent ICMP ping and traceroute messages from traversing the network.

Pings help to determine whether particular devices are reachable across an IP network and to measure network metrics, such as the time it takes for packets sent from the local device to reach a destination device, as shown in Fig. 5. Pings also measure the reliability and quality of a route by

monitoring packet loss. During network startup or commissioning and testing, the ping utility is a very valuable tool to help identify simple network connection failures.

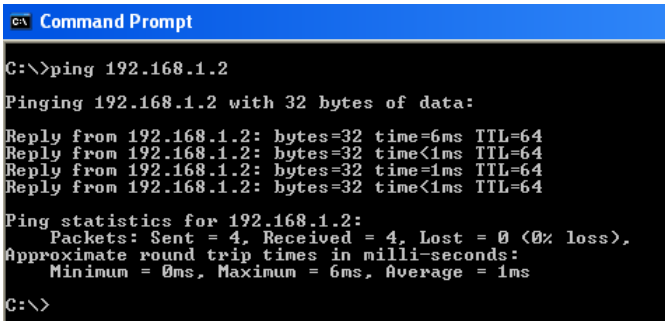

Fig. 5. Example of a ping from the DOS command prompt

Unfortunately, attackers can use the ping utility to gain valuable information about a network topology. A technique known as a ping sweep identifies devices that are active on a network. This knowledge can be used to craft an attack specific to a given network.

The ports and services that are active on a network are of primary concern when configuring a firewall. Port numbers fall into three ranges:

- Ports 0 to 1023, often called the well-known port numbers, are reserved for networking services that nearly all networked devices support.
- Ports 1024 to 49151 are the registered port numbers and are assigned to specific applications that networked devices may or may not support.
- Ports 49152 to 65535 are the dynamic port numbers. These ports will be opened or closed by a networked device as needed, usually in response to needs by certain applications, such as FTP.

Port 23 (associated with TCP) is the port assigned to Telnet. Telnet is a communications service for remote terminal sessions. It provides engineering access into the substation, even allowing access to an intelligent electronic device (IED). The service transmits all data in cleartext, including authentication data. This provides an opportunity for eavesdroppers to "sniff" the data, collect the data, and use the data at a later time for malicious purposes. It is important to configure the substation firewall to drop Port 23 packets to mitigate this type of attack, unless Telnet is absolutely necessary.

Secure Shell (SSH) is a preferred alternative to Telnet because it encrypts all transmissions. If SSH is not available, routing Telnet traffic through an Internet Protocol Security (IPsec) tunnel mitigates the risks of using Telnet.

Port 79 is the finger port. It identifies users of network devices. The information this service shares includes usernames, user email addresses, and the contents of the *.project* and *.plan* files located in the user's home directory. The availability of this information was not a concern in the early days of networking. It is more of a concern now as attackers find this information useful in social engineering attacks. In some systems, a user can turn off this feature; if this is not possible, a firewall rule can block this port and service.

Ports 161 and 162 are SNMP ports. Hackers use this protocol to expose management data. Most of the toolsets for implementing SNMP offer discovery mechanisms, which allow for remote administration of networked devices. Blocking these ports prevents unauthorized access to information about devices and users.

One of the more confusing protocols to deal with when implementing a firewall is the port associated with FTP. The FTP process begins by initiating a session on Port 21 of the FTP server, as shown in Fig. 6. Port 21 is used for session establishment and session control. The actual data transfer occurs on an available dynamic port. To transfer data, the FTP client and server establish a new session explicitly used for data transfer. When data transfer is complete, the data transfer session is closed. When a stateless firewall needs to pass FTP data, all of the dynamic ports must be left open to allow the session to switch to a high-level port. Otherwise, FTP data cannot be transferred. Stateful firewalls monitor the state of the connection and create dynamic rules to allow FTP and other dynamic protocols to switch ports. Like all cleartext protocols, it is better to operate an FTP session within the protection of a virtual private network (VPN) tunnel. Secure alternatives to FTP are Secure File Transfer Protocol (SFTP) and Secure Copy Protocol (SCP). A stateful firewall follows the changes in the state and port number and therefore protects the data from these types of connections.

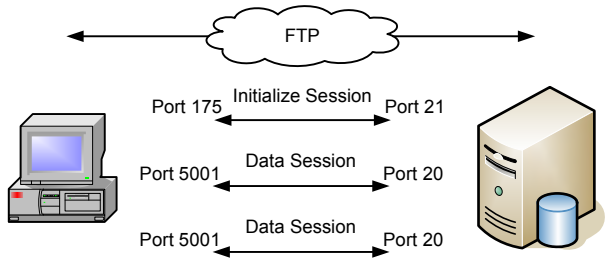

Fig. 6. FTP port numbering

Of primary concern to control systems are the DNP3 and Modbus® protocols. Firewall rules need to be configured to allow Port 20000 for DNP3 activity. The Modbus TCP port is 502, while the Modbus UDP port is 4800. The firewall rule allows for Modbus traffic to flow to and from the SCADA network out to the substation, as shown in Fig. 1.

V. Firewalls and Regulatory Requirements

Firewalls help to support North American Electric Reliability Corporation (NERC) Critical Infrastructure Protection (CIP) and National Institute of Standards and Technology (NIST) requirements, such as those found in NIST SP800-53 and SP800-82. The NERC CIP and NIST requirements and guidelines propose protecting critical

cyberassets, including protecting the electronic security perimeter (ESP), ensuring that critical cyberassets within an ESP gain protection from unauthorized access. The following are areas that a firewall supports, allowing firewalls to meet the NERC CIP requirements:

- Access control (CIP 003-2, R5)
- Electronic security perimeter (CIP 005-2, R1)
- Electronic access controls (CIP 005-2, R2)
- Ports and services (CIP 007-2, R2)

CIP 003-2 Requirement 5 states: "The Responsible Entity shall document and implement a program for managing access to protected Critical Cyber Asset information." [2]

Firewalls provide strong access controls and configuration information. These support compliance with this regulation.

CIP 005-2 Requirement 1 states: "The Responsible Entity shall ensure that every Critical Cyber Asset resides within an Electronic Security Perimeter. The Responsible Entity shall identify and document the Electronic Security Perimeter(s) and all access points to the perimeter(s)." [3]

Firewalls can define their configuration and the ESP, and log files assist in documentation of the access points to the substation network.

CIP 005-2 Requirement 2 states: "The Responsible Entity shall implement and document the organizational processes and technical and procedural mechanisms for control of electronic access at all electronic access points to the Electronic Security Perimeter(s)." [3]

A firewall implements strong access controls via rule sets that allow or reject traffic into a secured network.

CIP 007-2 Requirement 2 states: "The Responsible Entity shall establish, document and implement a process to ensure that only those ports and services required for normal and emergency operations are enabled." [4]

A stateful, deny-by-default firewall requires specific rules to open any port or service. Therefore, it operates as a compensating measure for products that cannot disable ports and services.

The NIST SP800-53 "Recommended Security Controls for Federal Information Systems and Organizations" provides guidelines for firewalls and calls out these devices to act as a security control or measure:

> Managed interfaces employing boundary protection devices include, for example, proxies, gateways, routers, firewalls, guards, or encrypted tunnels arranged in an effective security architecture (e.g., routers protecting firewalls and application gateways residing on a protected subnetwork commonly referred to as a demilitarized zone or DMZ). [5]

The NIST SP800-82 "Guide to Industrial Control Systems (ICS) Security" also suggests firewalls as a means for the following:

> Restricting logical access to the ICS network and network activity. This includes using a demilitarized zone (DMZ) network architecture with firewalls to prevent network traffic from passing directly between the corporate and ICS networks, and having separate authentication mechanisms and credentials for users of the corporate and ICS networks. [6]

A DMZ is a location on a corporate network that allows untrusted access to company data or information, such as what products and services a company has to offer or what numbers to call in case of a power outage. The types of services available might include web and email servers. A secure network design for a DMZ is to use two firewalls. One firewall that creates a DMZ, called the "front-end" firewall (shown in Fig. 7), is set up with rules to only allow traffic to the appropriate ports and services found in the DMZ. The second firewall, sometimes referred to as a "back-end" firewall, allows only specific or designated traffic from the DMZ to the company internal corporate network. Having layered firewalls makes an attacker work twice as hard to gain access to the internal network where the sensitive data reside.

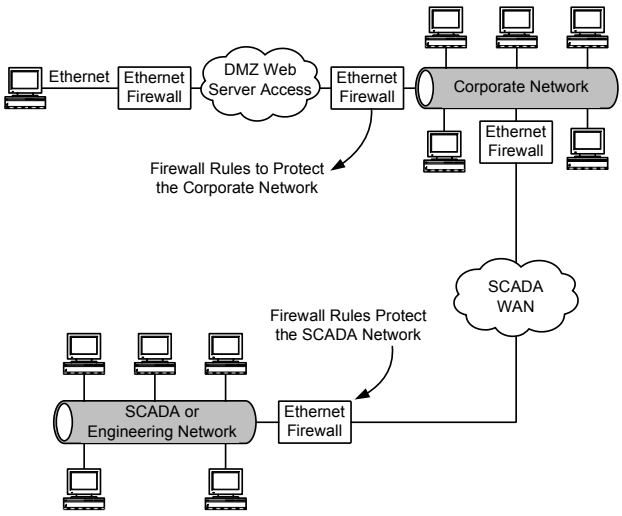

Fig. 7. Typical network topology with DMZ and firewalls

VI. Administration of Firewalls

Access to firewalls for administrative purposes is just as important to secure as the internetwork communications that firewalls protect. The primary concerns to address are remote management access to firewalls and securing this access. HyperText Transfer Protocol Secure (HTTPS) is one solution to secure remote access. HTTPS is only available on firewalls that utilize or contain a web server with a Secure Sockets Layer (SSL) or Transport Layer Security (TLS) protocol. SSH is another option that some firewalls use for secure remote access.

Although firewalls are a very useful tool for securing network communications, configuration problems can range in severity and scope. For example, rule changes to a firewall can inadvertently block authorized users from gaining access to network resources. Users can mitigate the chances of this happening by implementing a well-defined process and then verifying configuration changes prior to implementing them. For auditing purposes, firewalls should document all configuration changes. The firewall should document the change with the username of the entity that performed the change and the time and date the change occurred.

Firewalls that implement user-based accounts increase the authentication, authorization, and accountability of the system. User-based accounts authenticate a user to an identity rather than a role. This means that the firewall knows the identity of the person performing any action on the system. Knowing the identity of a user makes event tracking and event forensics more accurate. User-based accounts also ease the burden of password management by letting users manage their own passwords individually.

Firewall logs are also very beneficial to system administration, user education, and event tracking. Firewall logs provide a means to determine who did what and when. Logs also correlate to changes in the behavior of the network communications. This correlation provides useful insight in identifying the firewall rules that are causing system communications errors.

How firewall administrative ports handle failed login attempts is another important consideration prior to implementation. It is important to log all login attempts and lock accounts after multiple failed login attempts. Multiple login attempts can indicate a password cracking attempt. Automated password cracking devices will repeatedly attempt to access an administrative port to crack the password in a brute force attack. By implementing an administrative function that locks out a user for a period of time after three to five attempts, a firewall interrupts automated attempts to break passwords.

Firmware upgrade procedures for all embedded devices, including firewalls, are a concern. Firmware upgrades for updating security patches are essential to improve the reliability, functionality, and security of embedded devices. It is important to choose a firewall that requires authorization before allowing a firmware upgrade. Firmware upgrades that do not require proper user or firmware authentication can lead to permanent denial of service attacks or backdoors that attackers can use at will. Additionally, firewalls should document users that perform upgrades and tag the event with the date and time. It is important that upgrades be vetted on development systems first. At a minimum, the firmware upgrade process should require strong password access controls.

For the highest level of protection, consider cryptographic protection of firmware with hashes and digital signatures. Digital signatures can authenticate the firmware as having come from the device manufacturer. Hashes provide a method to ensure that the firmware did not undergo modification after it left the manufacturer's control.

It is also important for firewalls and all networking devices to meet or exceed the rigors found in the modern substation. Firewall specifications should include operational specifications that cover vibration, electrical surges, electrostatic discharge, fast transients, and extreme temperatures. It is important for these devices to also meet or exceed IEEE 1613 (Class 2) and IEC 61850-3 standards for communications networking devices in electric power substations.

VII. Conclusion

Firewalls are powerful cybersecurity tools that are found in many computer operating systems, antivirus software, and standalone products. Firewalls define the ESPs of computer networks and, if properly configured, prevent both attacks from accessing the network and the leakage of sensitive information. Layering firewalls between networks of different sensitivity levels increases the time it takes for an attacker to gain access to valuable information.

VIII. References

[1] Warriors of the Net. Available: http://www.warriorsofthe.net/.

[2] *Cyber Security—Security Management Controls*, NERC Standard CIP-003-2, May 2009.

[3] *Cyber Security—Electronic Security Perimeter(s)*, NERC Standard CIP-005-2, May 2009.

[4] *Cyber Security—Systems Security Management*, NERC Standard CIP-007-2, May 2009.

[5] *Recommend Security Controls for Federal Information Systems and Organizations*, NIST Special Publications 800-53, August 2009, p. 183.

[6] *Guide to Industrial Control Systems (ICS) Security*, NIST Special Publications 800-82, September 2008, p. 10.

IX. Further Reading

S. Northcutt, L. Zeltser, S. Winters, K. Kent Frederick, and R. W. Ritchy, *Inside Network Perimeter Security: The Definitive Guide to Firewalls, VPNs, Routers, and Intrusion Detection Systems*, 1st ed., Indianapolis, IN: New Riders Publishing, 2003.

H. F. Tipton and K. Henry, *Official (ISC)²® Guide to the CISSP® CBK®*, Boca Raton, FL: Auerbach Publications, 2007.

X. BIOGRAPHIES

Dwight Anderson received his B.S. in electrical engineering from Steven's Institute of Technology. He is now the security product manager for Schweitzer Engineering Laboratories, Inc. (SEL) in Pullman, Washington. Prior to joining SEL in 2005, he worked 20 years for Hewlett-Packard as an aerospace and defense business development manager and systems engineer, working on projects ranging from electronic warfare countermeasures to SCADA system programming. He is an active member of the FBI InfraGard forum, regarding the exchange of information related to critical infrastructure protection, and ISSA. He holds the Global Security Essentials Certification (GSEC) from Global Information Assurance Certification (GIAC) and is a Certified Information Systems Security Professional (CISSP).

Nathan Kipp received his B.S. in computer engineering in 2004 and his M.B.A. in 2007 from Washington State University (WSU). He joined Schweitzer Engineering Laboratories, Inc. (SEL) in Pullman, Washington in 2008 as a security applications engineer. Prior to joining SEL, Nathan worked as a systems administrator for the WSU Libraries. He participates in many industry organizations and is a member of ISSA and an Associate of $(ISC)^2$.