Boundary Protection for Transmission Facilities

Jeff Pack
POWER Engineers, Inc.

Presented at the CIGRE US National Committee
2018 Grid of the Future Symposium

SUMMARY

This paper discusses the requirements and cybersecurity architecture and controls associated with protecting the network boundary for transmission facilities. Most of these cybersecurity controls are commonly deployed, but boundary protection is still the leading issue identified by the National Cybersecurity and Communications Integration Center's (NCCIC) Industrial Control Systems Cyber Emergency Readiness Team (ICS-CERT) [1]. This paper explores why boundary protection is difficult to perform correctly and consistently, as well as areas to check when assessing a boundary protection system. Several examples of boundary protection devices and configurations are presented, analysed and updated to address issues.

KEYWORDS

Cybersecurity, network, boundary, firewall, authentication, authorization, logging.

jeff.pack@powereng.com

INTRODUCTION

One of the fundamental security controls for electric power transmission facilities is the design and deployment of boundary protection. The North American Electric Reliability Corporation (NERC) Critical Infrastructure Protection (CIP) Standard for defining an Electronic Security Perimeter (ESP) is detailed in NERC CIP-005-5 [2]. This requirement applies to transmission facilities with High or Medium Impact Bulk Electric System (BES) Cyber Systems and their associated Protected Cyber Assets with routable protocol connectivity.

In the Guidelines and Technical Basis for NERC CIP-005-5, NERC provides the functions for an ESP:

- Defines a zone of protection around the BES Cyber System.
- Provides clarity for entities to determine what systems or Cyber Assets are in scope and what requirements they must meet.
- Defines the scope of Associated Protected Cyber Assets that must also meet certain CIP requirements.
- Defines the boundary in which all Cyber Assets must meet the requirements of the highest impact BES Cyber System that is in the zone (the high-water mark).

Secure network design has many historical references to defining a boundary. In 1994, Bellovin and Cheswick [3] define security domains as a set of machines under common administrative control, with a common security policy and security level. That definition still holds true today.

With lots of history and millions of boundary protection devices installed worldwide, why is boundary protection the leading issue identified by NCCIC's ICS-CERT when performing assessments? There are some indicators in the FY 2016 report:

- Inadequate boundary protection for industrial control system (ICS) networks.
- No logical separation of the ICS from enterprise networks or untrusted networks (such as the Internet).
- No dedicated jump server to provide access to ICS data.
- Using the same user authentication credentials for the jump server and the enterprise network.
- Too many communication flows (ports and services) allowed to the boundary devices.
- Inadequate security services available in the demilitarized zone (DMZ) to support ICS system patching and updates.

This paper explores these indicators and determines what steps should be taken to mitigate the associated risk.

Inadequate Boundary Protection

The first indicator demonstrates that in many cases, the boundary protection is not configured properly and allows too much access from other security domains such as the enterprise network, and potentially Internet-connected devices.

The first area to review is the number, location and type of access points that exist on the substation network. This access point configuration is dependent on the types of communication and protocols as well as operational requirements. For example, transmission facilities often have a primary and secondary communications paths for SCADA information and in some cases, they may have to provide third-party access to a subset of data points. A simple example of this potential configuration is shown in Figure 1.

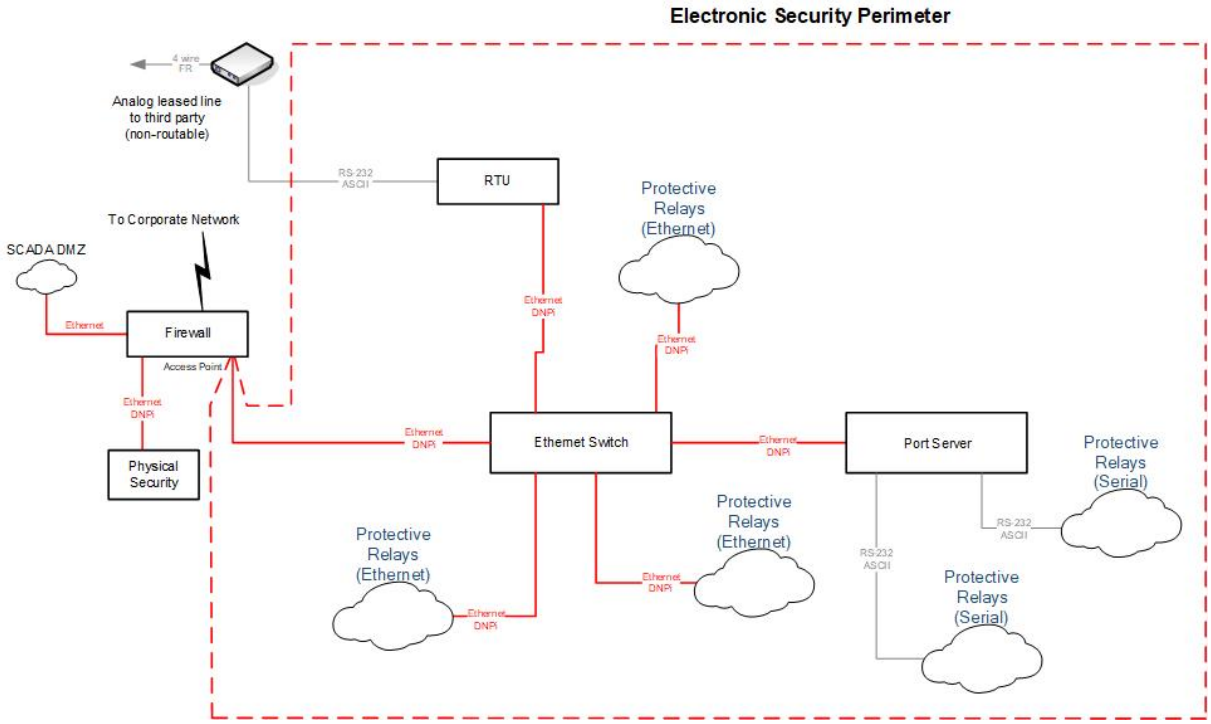

Figure 1 – Simple Substation Network Diagram

In this example, we have two potential access points identified – the firewall Ethernet port and the serial port on the Remote Terminal Unit (RTU). There are several areas where additional information is required to determine the true level of risk from this configuration:

- Firewall configuration
- Modem and RTU configuration
- Physical access

The firewall requires proper configuration to protect the substation network from unauthorized access. The firewall should be configured to only allow access from a limited number of pre-defined devices from the corporate network; ideally only from a jump server that provides additional authentication and encryption of remote access traffic.

Ensure that all firewall changes are reviewed and approved through a configuration management process. The failure to track all changes to a system is a major factor in security systems becoming less effective over time [4].

The firewall should also only allow the required network ports and services to perform the operational requirements of the substation. For example, if there is no operational requirement to allow HTTP or HTTPS access to the equipment on the substation network, then those services should be blocked by the firewall configuration.

Firewalls themselves often require communication between a management console and the individual firewall devices at remote locations. It's important to understand what services are required between firewall devices for operation and confirm that those services are required. Firewalls may have a separate set of implied rules that are enabled by default and are not normally displayed on a management console view of the active firewall rules or on an operator's console. Review your firewall configuration carefully for terms such as:

- Accept control connections
- Accept update connections

- Accept dynamic routing updates
- Accept Domain Name Service
- Accept Web connections for firewall administration

A comprehensive understanding of the management architecture of the firewall systems is essential in maintaining a secure configuration across all access points. Some firewalls have global parameters that are applied to all devices under management by a specific console or instance, while other implementations define parameters that require customization for each individual device or software feature enabled. The configuration management process provides a check and balance to ensure that firewall devices remain configured properly.

Verify that your firewall design and maintenance processes actually require all the enabled services and disable any services not required.

The modem and RTU configuration should require authentication to establish the dial-up connection. While the connection may appear to be only SCADA data, the configuration on the RTU should be reviewed to determine if remote access is possible from the dial-up connection. Some RTU and serial protocols allow both data and control commands to be carried at the same time.

Finally, physical access controls to the substation and control network are a major security design requirement – otherwise a knowledgeable attacker can quickly access many of the devices or install their own remote access device to exploit the network later from a safe distance.

No Logical Separation of ICS Network

The second indicator discusses the lack of logical separation of the ICS network from corporate or other networks or even the Internet. While this may sound like the previous indicator, there are several elements included in logical separation of networks that go beyond the firewall.

Logical network separation can involve several different design elements at multiple locations in the network. With new technology enabling virtualization at nearly every level of the traditional layered network model, it's important to review logical separation at each layer.

In an Ethernet network, logical separation begins at the data link layer. As indicated in the Guide to Industrial Control Systems Security developed by the National Institutes of Standards and Technology (NIST) SP 800-82 R2, VLANs can provide logical separation and segregate traffic for ICS networks at layer two [5]. ICS networks should not use the default VLAN in their design, as many devices added to a network will default to VLAN 1 and therefore may become active on the VLAN.

In many modern substation networks, the Ethernet network is used to carry IEC 61850 protocol traffic which may include time-sensitive protection and control functions using Generic Object Oriented Substation Event (GOOSE), so the proper design and separation of virtual Local Area Networks (VLANs) is critical for proper operation. VLANs should be grouped by operational function and all devices participating in GOOSE messaging associated with a specific function should belong to the same VLAN and class of service priority to avoid latency issues [6].

Most network switches allow for the definition of port and VLAN access control lists, or a forbidden port list for VLAN membership. Use of these mechanisms to limit VLAN membership by physical ports are recommended to minimize the chance of device misconfiguration and improper traffic on VLANs. As with the firewall management tasks described earlier, a robust configuration management process helps ensure that network switches are configured properly.

Moving up to the network level, we introduce the ability to route traffic between local area networks. This is the traditional location where firewalls and routers separate network segments into security

domains. The ability to route network traffic is an important distinction in the NERC CIP Standards and underscores the additional risks associated with routable network protocols.

Some of the fundamental architectural and design choices for boundary protection at layer three include:

- Deny all network traffic by default and allow network connections by exception to ensure that only approved connections are allowed.
- Defining a demilitarized zone between the ICS security domain and all other security domains. This architecture is covered in more detail in the Inadequate Security Services to Support ICS System Maintenance section of this document.
- Only allowing communication between specific, pre-defined source and destination addresses that require approval and periodic review.
- Disabling control and troubleshooting services and protocols such as SNMP and DNS.
- Configuring boundary devices such as firewalls and routers to fail in a predetermined state based on the operational priorities of the organization.
- Configuring ICS security domain devices with separate network addresses that are fundamentally different than corporate network addressing.

At the transport and application layers, the focus should be on defense-in-depth and providing a secure configuration of operating system and application software. Since the focus for this paper is on boundary protection, there are no specific details due to the wide number of operating systems and application software packages that may be installed in the ICS security domain. However, there are some architectural elements that can be applied across the transport and application layers:

- Develop and enforce an effective set of policies, procedures and processes that describe and govern the security program for the ICS security domain.
- Implement the principle of least privilege throughout the overall ICS system. Only configure the communications and application access and features required to achieve the operational requirements of the system. For example, there is normally no need to allow operator accounts to make configuration changes, or to allow protective relays to communicate directly with the revenue meters; therefore, there should be an access control list or a firewall rule that prevents that type of communication from happening. As a corresponding issue, if that type of access or traffic is detected on the network, the reason should be investigated immediately as malicious code may be trying to configure or communicate to any device that will respond.
- Implement and enforce configuration and change management processes to provide a review and approval element for all hardware and software configurations and help maintain an acceptable risk level.
- Implement an aggressive patch and vulnerability management program to assess, detect and patch vulnerabilities across all elements of the ICS security domain.
- Implement an extensive security monitoring system including network security monitoring, centralized log management and security incident and event monitoring.

No Dedicated Jump Server for ICS Access

In many cases, firewall rules are defined to allow maximum flexibility for staff to access the ICS security domain from any remote location and connect to devices to provide troubleshooting and other maintenance functions. However, one of the major design goals for NERC CIP was to minimize the risk of having routable protocols available outside the ESP.

NERC CIP-005-5 describes the implementation of an intermediate system or jump server that resides between the ICS security domain and the corporate network and serves as a proxy for a remote user who wants to access devices in the ICS security domain. This eliminates any direct connection between external devices and all the devices in the ICS security domain.

There are several benefits to this architecture for interactive remote access:

- Firewall rules can be much more restrictive for inbound traffic to the ICS security domain – all remote access is forced to use the jump server as the source of the network connection.
- The jump server can enforce multi-factor authentication prior to any connection attempts to devices in the ICS security domain.
- The jump server can be configured to use an encrypted connection between the remote computer and the jump server to protect authentication credentials and other sensitive information.

Based on these benefits, the use of a jump server for interactive remote access is required for High Impact and Medium Impact BES Cyber Systems and their associated Protected Cyber Assets with external routable connectivity.

Using Common Authentication Credentials

This indicator refers to the use of the same username and passwords on the corporate network and the jump server. One highly visible example of this issue is the Ukrainian Power Grid Attack of December 2015. In its analysis [7] of this event, the Electricity Information Sharing and Analysis Center (E-ISAC) reviewed the attack techniques used to gain access – using malware to harvest credentials to the corporate network and then pivoting their efforts to access the control network where the SCADA dispatch workstations and servers existed.

While multi-factor authentication is a good practice, if the same username is used in both authentication domains, the attackers have a head start in gaining access to the jump server. If the authentication domains are managed so that there are no duplicates in usernames, the overall security posture of the jump server, and in turn the ICS security domain, is increased.

Any administrative changes to the authentication and authorization systems should be reviewed and approved per the configuration management process due to the potential impact of a failed authentication system. In addition, use of a centralized authentication service such as LDAP, Active Directory or RADIUS allows for easier credential maintenance and easier automation of disabled and removed user accounts. Some jump servers also provide the ability to manage default or other generic accounts on intelligent electronic devices (IEDs) such as protective relays, fault recorders and meters. Depending on the number of IEDs in the substation, this may be a valuable feature since NERC CIP requires password changes for these accounts every fifteen calendar months.

Excessive Network Services to Boundary Devices

This indicator is related to the need for logical separation of the ICS security domain, but is more focused on minimizing access to the servers housed in the DMZ and the access points for each ICS security domain. In many cases, these servers are treated similarly to servers on the corporate network that provide many different services to a broad category of users. Corporate servers are normally allowed access by default rather than by exception since nearly all employees use those servers daily. Servers that provide dedicated support to the ICS security domain should only be accessed by those devices that truly need access and only using the minimum number of services required by the operational requirements.

For example, a jump server located in the DMZ should only be accessed by approved users with ICS security domain accounts and only from certain locations on the corporate network. There should be no reason for the finance department to access the jump server, so those department accounts and IP addresses should be restricted from access.

Inadequate Security Services to Support ICS System Maintenance

A common practice in ICS system design is to allow outbound only access to hardware and software vendors for patching and software updates, including anti-malware signatures. While this sounds reasonable to support these maintenance activities, it also allows communications directly from ICS devices to insecure networks and offers attackers a way to access the ICS security domain.

A better solution is to provide these services on the DMZ. The architecture is depicted as indicated in the Guide to Industrial Control Systems Security (NIST SP 800-82 R2) below. The devices in the ICS security domain would access one or more servers located in the DMZ to support hardware and software patching, anti-malware signature updates and other data exchange services such as a data historian. The corporate network can access the data historian and provide updates to the other servers in the DMZ but would not be allowed to directly access any server in the ICS security domain.

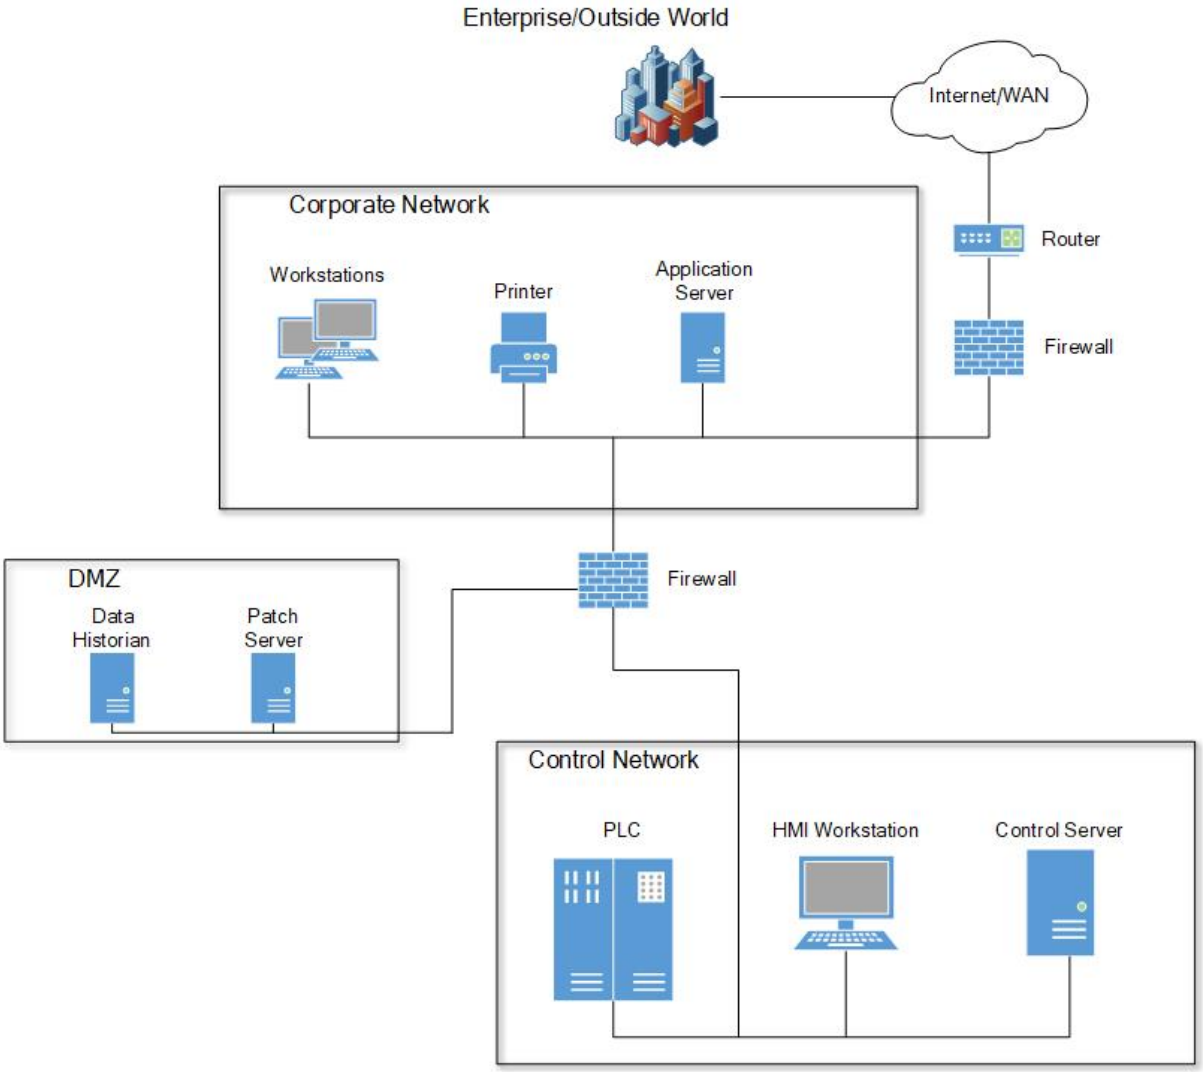

Figure 2 – DMZ Network to support Control Network (from NIST SP 800-82 R2)

Summary

The technology to provide effective boundary protection for ICS networks is not new or even very sophisticated compared to the complex control systems that are required to support the future power grid. Boundary protection provides operational challenges and limits some of the opportunities for convenient remote access and data exchange. Each organization will need to design a solution that meets its risk management expectations and strikes the appropriate balance between security and operational function.

As ICS-CERT indicates, however, it is difficult to implement and maintain an effective boundary protection system. This paper provides a review of identified boundary protection issues and suggested actions to incorporate into the boundary protection architecture, design and implementation. While there are several issues raised and recommended steps to mitigate associated risk, the hardest part of maintaining an appropriate cybersecurity risk profile may be the need to maintain, track and approve all configuration changes to the systems.

Organizations that implement a sound architectural design and maintain good configuration management of boundary protection will have an effective operational solution that maintains the desired risk profile.

BIBLIOGRAPHY

[1] NCCIC, ICS-CERT Annual Assessment Report, FY-2016. https://ics-cert.us-cert.gov/sites/default/files/Annual_Reports/FY2016_Industrial_Control_Systems_Assessment_Summary_Report_S508C.pdf

[2] NERC, CIP-005-5, Cyber Security – Electronic Security Perimeter(s). http://www.nerc.com/pa/Stand/Reliability%20Standards/CIP-005-5.pdf

[3] Bellovin, Steven and Cheswick, William, "Firewalls and Internet Security," Addison-Wesley, 1994.

[4] Shackleford, Dave, "Secure Configuration Management Demystified," SANS Reading Room, August 2012. https://www.sans.org/reading-room/whitepapers/analyst/secure-configuration-management-demystified-35205

[5] NIST Special Publication 800-82, Revision 2, "Guide to Industrial Control Systems (ICS) Security," May 2015. https://doi.org/10.6028/NIST.SP.800-82r2

[6] Mulholland, Daniel and Strydon, Gideon, "Automation Solutions Using IEC61850 GOOSE," PACWorld, June 2015. https://www.pacw.org/issue/june_2015_issue/lessons_learned/automation_solutions_using_iec_51850_goose.html

[7] E-ISAC Defense Use Case, "Analysis of the Cyber Attack on the Ukranian Power Grid," March 2016. https://ics.sans.org/media/E-ISAC_SANS_Ukraine_DUC_5.pdf