CIGRE-425

2021 CIGRE Canada Conference & Expo
Toronto, Ontario, October 25-28, 2021

Securing a substation LAN with a Software Defined Network and a Functional Security Monitoring Intrusion Detection system

Steel McCreery
OMICRON electronics Canada Corp
Canada

Lance Dice
Schweitzer Engineering Laboratories Inc.
USA

steel.mccreery@omicronelectronics.com
lance_dice@selinc.com

SUMMARY

Since the first confirmed blackout caused by hackers in the Ukraine in 2016, the frequency and intensity of cyber-attacks targeting the power grid have increased. Substations represent a critical attack vector within the power grid and these substations have unique challenges that are not found within traditional IT or control centre networks. Commonly located in remote locations, unmanned, and visited infrequently by a transient workforce increasingly comprised of contractors, the probability of potential introduction of viruses and other malware is high. The more sophisticated malware is designed specifically to perform tasks such as network and traffic reconnaissance in order to lay the groundwork for more sophisticated attacks. Gateways and firewalls, traditionally located at the substation LAN perimeter, cannot prevent the transmission of unauthorized network traffic within the cyber defence perimeter. During these intrusions, the generation of additional network traffic or a slight change in behavior of a device may be the only indicators of the presence of an intrusion or malware infection. This paper examines the benefits of securing substation LANs thorough the complementary security features of Software-Defined Networking (SDN) and a functional security monitoring Intrusion Detection System (IDS). The paper first examines common attack vectors and exploits found within the substation LAN and devices. The paper examines how the SDN configuration process derives the correct communications flows for all devices within the LAN. These flows determine the paths that each frame of ethernet traffic (for each conversation) must take from the source to destination device. The paper then explores how each frame is examined from layer one through four at each interposing SDN switch to ensure it contains the correct match criteria before it can egress the SDN switch. The frame will then go on to the next SDN switch within the path and will have its match criteria examined yet again. The SDN LAN can be configured to automatically send a copy of every ethernet frame from every communication to the Intrusion Detection System (IDS). The inherent shortfalls of the traditional signature-based and learning-based intrusion detection systems when applied to substation LANs are examined before an explanation of the configuration and operation of the functional security monitoring IDS is presented. The functional security monitoring IDS is tailormade for the substation. In the case of an IEC 61850 substation, this form of IDS uses the information contained within the SCD file (which includes information describing the entire automation system, the devices, their data models, their communication patterns, the primary assets and potentially the single-line diagram of the substation) to build a precise system model of the automation and power system communications. This precise model allows the IDS to compare each packet received from the SDN network against this live system model to measure performance characteristics in order to detect hardware failures or malfunctions in addition to the ability to perform deep packet inspection. Deep packet inspection is the process of comparing the variables contained within each communications service (GOOSE, MMS and SV) against the expected content derived from the system model. By using the substation section within the file, a single line diagram "like" overview can be created that allows for graphical representation of events in addition to textual alarm messages that are easily comprehended. Further, the functional communications security monitoring IDS can assign functions to non-IEC-61850-devices such as a “testing PC” which determines what traffic the device can generate without causing an alarm.

KEYWORDS

IEC 61850, Software Defined Networking, Cyber Security, Functional Security Modeling Intrusion Detection System

INTRODUCTION

For this paper we will define a cyber-attack on a substation as an event where an attacker modifies, degrades, or disables a service of at least one protection, automation, or control device within the substation. Since the first confirmed blackout caused by hackers in the Ukraine in 2016, the frequency and intensity of cyber-attacks targeting the power grid have increased. Substations represent a critical attack vector within the power grid and these substations have unique challenges that are not found within traditional IT or control centre networks. Commonly located in remote locations, unmanned, and visited infrequently by a transient workforce increasingly comprised of contractors, the probability of potential introduction of viruses and other malware is high. The more sophisticated malware is designed specifically to perform tasks such as network and traffic reconnaissance in order to lay the groundwork for more sophisticated attacks. Gateways and firewalls, traditionally located at the substation LAN perimeter, cannot prevent the transmission of unauthorized network traffic within the cyber defence perimeter. During these intrusions, the generation of additional network traffic or a slight change in behavior of a device may be the only indicators of the presence of an intrusion or malware infection. This paper begins with a very brief overview of two foundational principles of cyber security; defense in depth and the NIST framework. From there, the attack vectors of a substation and counter measures for communication connections to the substation are discussed prior to examining the benefits of securing substation LANs behind the firewalls through the complementary security features of SDN and a functional security monitoring IDS.

One cybersecurity principle well accepted across all regulations worldwide is the defense-in-depth principle. This principle recommends that there be not one hard shell of security around your substation but multiple layers of defense like the layers of an onion. If one layer of defense is breached not all is lost, the attacker faces yet another layer of defense. A second well accepted principle is the principle of deny-by-default: focus on what should be there and deny everything else. In practice the physical barriers, fences and locks around the substation perimeter represent physical layers of defense while network segmentation, firewalling and role-based access controls are examples of the layering of cyber defenses employing the deny-by-default principle within the communications network. Working in concert with the defense-in-depth and deny-by-default principles is the NIST Cybersecurity Framework, which is a set of guidelines for private sector companies to follow to be better prepared in identifying, detecting, and responding to cyber-attacks. This framework is not only used in North America but other countries as well. The core assumption of this security framework is that there is no 100% protection- attacks can always come through the various layers of defense. Within the NIST Framework security is seen as a process that has five steps: *identify* assets and attack vectors, *protect* against the vectors with the highest risk, *detect* attacks/threats as they occur and *respond* to detected threats to minimize damage and learn.

Figure 1 depicts the substation attack vectors. Most attacks have used the remote access link to the office IT to gain access to the substation LAN. The Triton cyber-attack on plant safety PLCs and the Ukraine 2016 attack on a substation in Kiev are examples of previous attacks that have used this attack vector. Countermeasure for this attack vector include the creation of a DMZ: a separate substation IT network with its own servers and services. The services would consist of virtual machines to perform the tasks traditionally accomplished with maintenance PC's. These virtual machines would be connected to the substation LAN and would be accessible through a secure limited access link from the office

IT. The access would use multifactor authentication; a token device (the engineer would need a code from the token) and a password to gain access to the virtual machines. This would effectively eliminate impersonation and automated attacks. Firewalls on either side of the substation RTU are very effective countermeasure to secure this attack vector (which was used in the 2015 attack on a Ukrainian substation) other attack vectors beyond the firewalls still remain (Figure 2). They consist primarily of transient cyber assets used by maintenance and engineering personnel. Policies and procedures that address how transient cyber assets should be used within substations given no assurance these policies and procedures will not be inadvertently violated potentially resulting in the introduction of malware onto the network. The question then arises as to how potentially malicious traffic can be detected beyond the firewalls.

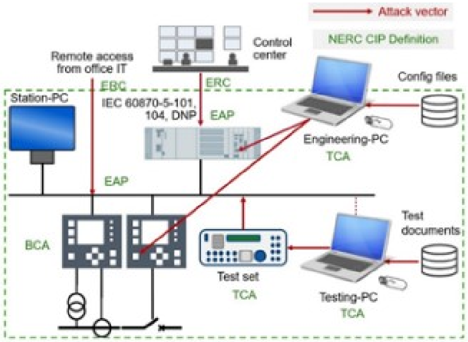

Figure 1

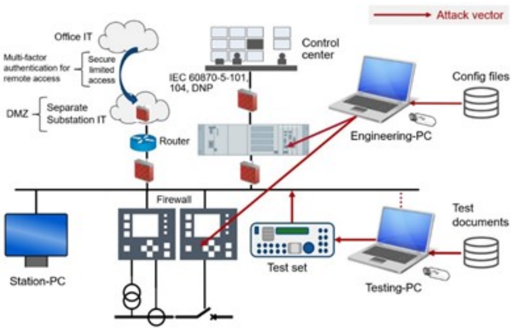

Figure 2

Traditional layer 2 switches are comprised of both a control plane and the data (forwarding) plane. The control plane examines the ethernet header portion of the frame to determine how to forward the frame. Standard layer two switch forwarding is based on trust, in that all frames are forward based solely on the content of the ethernet header and standard switch operating principles and there is no verification process to determine the validity of the frame or traffic that is being forward. SDN addresses this weakness directly: SDN is an abstraction that moves the control from networking devices to a centralized controller by decoupling the switch control plane from the data plane. Within SDN the term "flow" refers to data flow from a source to destination that is permitted based on characteristics of the frames and rules stored within the "flow table" of each interposing SDN switch. Given that SDN forms the LAN which moves all traffic, there are no blind spots: the SDN sees all traffic on the LAN. The SDN switches examine Ethernet frame data in OSI layers 2 through 4 (figure 3), such as MAC addresses, EtherType, VLAN, IP addresses, and TCP/UDP ports in addition to the physical ingress port and applies rules stored within the switch flow tables to determine an associated action to take for each frame that ingress the switch. Automated tools allow the network engineer to configure the source and destination(s) switch port(s) for each flow and the associated rules. At the time of flow configuration, the network engineer has the option to have a failover path (an alternate path through the SDN network) automatically configured, such that in the event of a primary path component failure, the frame has an alternate path to its destination which provides sub-millisecond failover times and eliminates the need for spanning-tree protocols. Once the flow is configured, frames that ingress the SDN network at the specified switch will have these flows rules applied. If the frame satisfied the rules of a flow, the switch will follow the instructions in the flow. This process continues until the frame reaches the designation port(s) to egress from the SDN LAN. In effect, each switch is a rudimentary firewall checking and rechecking the validity of each frame as it traverses the LAN. The injection of a frame that does not meet preconfigured rules will be immediately detected at the SDN switch port it was injected into.

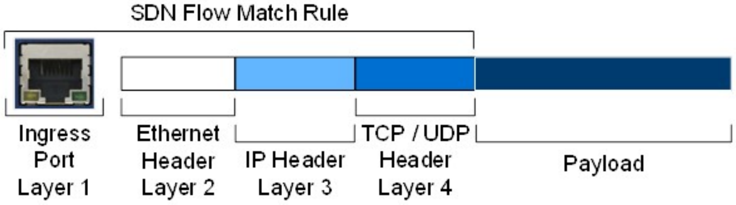

Figure 3

An SDN network can be categorized as either reactive or proactive. Reactive SDN operation is when a switch receives a packet that does not match any of the flow table entries and forwards the packet to the flow controller. The flow controller then decides how to handle the packet. It can drop the packet, or it can add a flow entry that tells the switch how to forward similar packets in the future. This type of SDN operation is common in IT networks where there is a great deal of variation in the type of traffic on the network. The problem with this mode of operation is that it requires the controller to always be active, turning the controller into a single point of failure. Unlike the IT LAN, the characteristics of all traffic is well known within an OT LAN, allowing the implementation of proactive SDN. During the proactive SDN configuration process, the engineer configures rules based on layer one through four to create the primary communication "flows" or permissions for all known traffic. In the case of an IEC61850 substation, the SCD file containing the communications configuration of all devices can be used by some SDN systems to automatically configure all substation flows, negating the requirement for port-based VLAN configuration. Even though VLANs are not configured per-port, the SDN switches will still only forward the messages to only the devices that need them through matching the VLAN on the packet and the destination MAC address. Once all flows have been created and the flow tables (or rules) have been downloaded to the switch's non-volatile memory, the controller's primary task is complete and it then performs auxiliary functions, such as the collection of network statistics and visualization of the network status for the user interface. If the controller were to fail or be removed from the LAN, the LAN would continue to run unaffected.

As stated earlier, a key step within the NIST framework is the detection of attacks or threats as they occur. The SDN LAN is the first layer of detection beyond the firewall and can be configured to send a copy of all traffic to egress the SDN switch port connected to the intrusion detection system. This reduces the cost and complexity of setting up mirrored ports on every traditional switch to forward a copy of all traffic to the IDS.

Traditionally there are two main categories of IDS: the signature-based IDS and the learning-based IDS. Signature-based IDS's rely on a blacklist built from prior knowledge of key characteristics or, if preferred, signatures of previous attacks. The IDS scan the data stream for known patterns from this list, similar to how a PC virus scanner scans the PC's memory and drives for virus signatures. The problem with this approach is that there have been only a few known attacks on substations and so there is limited data to discern a signature or blacklist from. Further, it is likely that a new attack will present a new signature not known by the IDS, having a high probability of going undetected. Given that the first occurrence of this new attack could have severe consequences, a substation IDS must be able to detect attacks without any previous knowledge of what the attack might look like, which is a very different approach then that employed by the signature-based IDS.

The learning, or so-called artificial-intelligence-based IDS, requires a learning phase of several weeks to learn the typical traffic patterns of a healthy system to develop rules. Such systems look at frequency and timing of certain protocol markers to attempt to learn the usual behavior of the system. Once this learning phase has been completed, the IDS monitors the traffic and an alarm will be raised if one of the markers is significantly outside the expected range. The main issue with this approach is that valid, but infrequent, traffic scenarios may not take place during the learning phase such as a real protection trip, uncommon switching or automation actions, or routine maintenance and testing traffic, all of which are valid network traffic but will result in a number of alarms. Another problem with this approach is the way in which such alarms are reported. These systems don't understand the semantics of the protocols, and so alarm messages are expressed in terms of technical protocol details. Hence, alarms can only be examined by an engineer skilled in IEC 61850 protocol details and familiar with IT network security. The engineer examining the alarm must also know the operational situation to judge if certain IEC 61850 protocol events correspond to valid behaviors. As a result, the decoding is inevitably labor intensive and difficult. Once the cause of the alarm is known, and the traffic to be found valid additional time and effort must be taken create rules to ensure the IDS no longer alarms on this valid traffic. Over time this effort proves to be a daunting task as there is a tendency to be a high number of false alarms and there are many substations all of which requiring the same highly skilled personnel to examine and root cause each and every alarm and if the traffic determined valid time must be taken to create rules to ensure the IDS does not alarm on the traffic. Given the effort to address each alarm, over time, this situation often leads to the alarms being ignored or alarms discarded without investigation and eventually the IDS being disabled altogether.

In this paper we will use IEC61850 services to explain capabilities of a Functional Security Monitoring IDS. As mentioned earlier, for IEC 61850 substations the entire automation system, including all devices, their data models, and their communication patterns are described in a standardized format – the System Configuration Language (SCL) file. The System Configuration Description (SCD) files can be generated by engineering configuration tools that can contain addition information about primary assets and potentially the information to create single-line diagram. Communications Modeling is a term that describes a new approach to the detection of intrusions, by deriving a precise model of the automation and power system communications from the information contained within the SCD file. This precise model allows the IDS to compare each packet received from the network against this model to verify the traffics validity. Further, the model allows the IDS to perform deep packet inspection: the IDS can compare the variables contained within the communicated (GOOSE, MMS, SV) messages against the expected content. This level of inspection is made possible without the need for a learning phase: the IDS configuration process is initiated by the import of the SCD file. The information contained in the SCD file allows the IDS to create the system model and a whitelist of valid communication messages. Through a comparison of the actual network traffic to that of the system model, zero (or minimum) false alarms are expected. The ability of the IDS to go deeper into the traffic to inspect the content of the message (GOOSE, MMS, or SV) and compare it to what is expected, brings the IDS monitoring to a new level of capability: the IDS is capable of what is termed deep packet inspection. This information is used by the IDS to perform functional security monitoring. We will use the information within a GOOSE message to better understand functional security modeling. In addition to the actual data, the GOOSE message contains additional codes that the IDS can draw from in order to detect functional errors. For example, monitoring the state number (stNum) and sequence number (sqNum) of each GOOSE message received allows the IDS to detect glitches (missing GOOSE messages) in the sequence of GOOSE messages being received, which can generate an alarm. This type of alarm is an indication of intermittent network communications or that a device has been restarting. Either situation is valuable information for the operator, indicating well the system is functioning. Let's look at another example: at the time of publication, each GOOSE message has a time stamp called the Entry Time Stamp. Through a comparison of a GOOSE message Entry Time Stamp to that of the current time when the GOOSE message is received at the IDS port, the IDS can detect if a time synchronization issue exists within the device that published the GOOSE message. Another powerful capability of a functional security monitoring IDS is that of being able to classify the function of addition devices on the LAN, such as RTU's, HMI's and Test PC's, and assigning "permissions": what communications these devices are permitted to generate without triggering the IDS to alarm. If for example a device is classified as a "Test PC" it would be given the "permissions" to generate traffic necessary for testing such as the traffic required to monitor the data model of an IED. The Test PC would not have the "permissions" to send a command to the IED to control a breaker. If the IDS detect such traffic from the Test PC the IDS would alarm. Further, given that the IDS understands the context of the alarm, it would have the ability to generate an accompanying message that would provide context for the alarm, making it easier to comprehend. The IDS can also utilize the information within the substation section of the SCL file to create an overview diagram of the substation (like a single line diagram), as shown in Figure 4, and superimpose an indication of what device or devices are associated with an alarm. Figure 4 shows an example of such an alarm: the vertical red arrow indicate which devices are associated with the alarm and the text message gives further context to the alarm. In this case the text reads "Switching command on AA1D1Q2Q1CONTROL/CSWI1.Pos", indicating the that Test PC tried to perform a switching operation on an IED. The associated text description of the alarm is clear and in the language of the PAC's engineer. Taking this example further, the Functional Modeling IDS could be configured with additional rules for the permissions such as a rule that if the IDS is informed the system is in a "Maintenance" or "Commissioning mode" the Test PC is allowed to generate commands to control the breaker without the IDS generating an alarm. These are just a few of the many added capabilities of a Functional Security monitoring Intrusion Detection System.

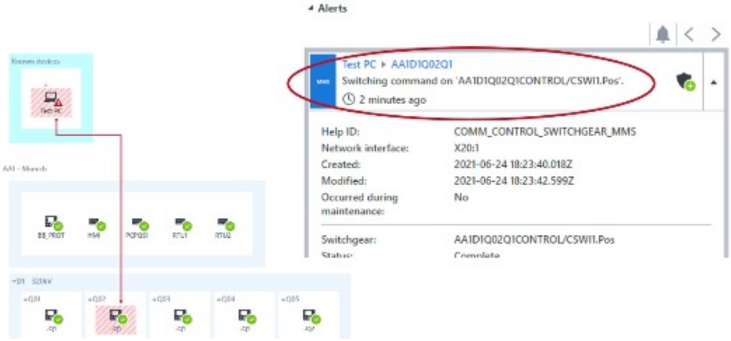

Alerts

Test PC ▸ AA1D1Q02Q1
Switching command on 'AA1D1Q02Q1CONTROL/CSWI1.Pos'.
2 minutes ago

Help ID: COMM_CONTROL_SWITCHGEAR_MMS
Network interface: X20:1
Created: 2021-06-24 18:23:40.018Z
Modified: 2021-06-24 18:23:42.599Z
Occurred during maintenance: No

Switchgear: AA1D1Q02Q1CONTROL/CSWI1.Pos
Status: Complete

Figure 4

CONCLUSION

Substations cyber-attack vectors must be determined and effectively protected to prevent potentially severe consequences for the grid. The complementary features of SDN and a Functional Security Monitoring IDS provide a very powerful and robust cybersecurity stance that is worth consideration.