Implementing Security for Critical Infrastructure Wide-Area Networks

Paul Robertson, Colin Gordon, and Simon Loo
Schweitzer Engineering Laboratories, Inc.

Published in
Sensible Cybersecurity for Power Systems: A Collection of Technical Papers Representing Modern Solutions, 2018

Originally presented at the
Power and Energy Automation Conference, March 2013

Implementing Security for Critical Infrastructure Wide-Area Networks

Paul Robertson, Colin Gordon, and Simon Loo, *Schweitzer Engineering Laboratories, Inc.*

***Abstract*—Cybersecurity is a major concern for individuals and organizations that manage and operate communications networks.**

Most modern utility and industrial processes rely on using high-bandwidth data communications networks to run and manage complex operations and systems. A typical electric utility communications network supports applications such as relay pilot protection, line current differential protection schemes, synchrophasor data collection, supervisory control and data acquisition (SCADA), engineering access, voice, surveillance, event report collection, and many other types of tasks. These same data communications systems can also provide opportunities for unauthorized access to these applications if appropriate cybersecurity measures are not established and implemented. The consequences of unauthorized user access can be costly and potentially catastrophic if networks associated with safety-critical systems are considered.

This paper describes several best practice approaches for securing wide-area network (WAN) communication for critical infrastructure applications.

I. Introduction

Most modern utility and industrial processes rely on using data communications networks to run and manage complex operations and systems. Unfortunately, there is an increasing trend in the number of cyberattacks against utility and industrial systems that involve exploiting weaknesses in the security of their communications networks. As a consequence, cybersecurity has become a major concern for individuals and organizations that manage and operate communications networks. This paper describes several best practice approaches, such as defense-in-depth strategies and strong user access controls, for securing wide-area network (WAN) communications for critical infrastructure applications.

II. Application of Wide-Area Networks for Critical Infrastructure

The purpose of a WAN is to transport data between distinct geographical sites. WANs typically carry large amounts of data for a wide variety of applications and services. Electric power systems, water treatment plants, large industrial facilities, petrochemical plants, gas pipelines, and transportation systems all rely on wide-area communication to run and manage their complex processes and systems. The majority of the processes in these operations are considered critical infrastructure, which is the term used to define infrastructure that is essential for the success and well-being of an economy [1].

WAN data communications types used by critical infrastructure devices can be divided into the following four general categories (in order of importance):

- Protection – required for immediate operations that are essential to the safety and reliability of critical infrastructure.
- Supervisory control and data acquisition (SCADA) – required for system visibility, metering, and automation applications.
- Engineering access – required for human-to-machine operations, including gathering reports and changing settings.
- Informational – required for other information technology (IT) applications, such as Voice over Internet Protocol (VoIP) and video data.

The substation, for example, is a critical infrastructure facility that may use all four of these traffic types. Fig. 1 shows a typical WAN for a power utility, with a separate edge network WAN providing communication to critical infrastructure substations and generation sites. The core network WAN is used to provide communication to corporate sites. A utility substation contains a diverse range of equipment that requires a WAN to support applications that include voice, teleprotection, video, control or automation, and possibly even corporate local-area network (LAN) access for other information functions.

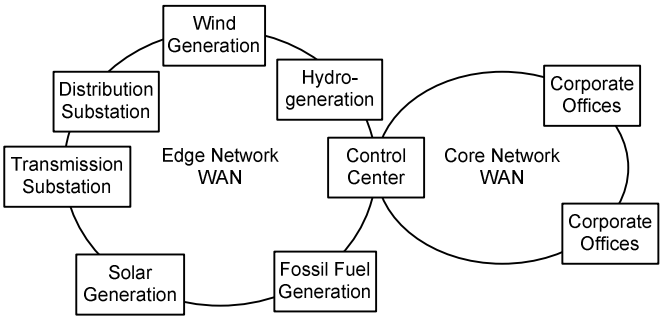

Fig. 1. Power utility WAN.

By its definition, WAN traffic ingresses and egresses controlled digital and physical local-area environments of an office, plant, or substation facility and, in some cases, interconnects with other WANs operated by telecommunications providers or even other critical asset owners. These WAN attributes introduce specific cybersecurity threats that expose the processes and systems within the critical infrastructure facilities to a greater

possibility of cyberattacks. It is therefore essential that cybersecurity threats to WANs be understood and appropriate steps taken to minimize the risk of adverse cyber-events.

III. Cybersecurity Threats to Critical Infrastructure WANs

For the sake of simplicity, this paper divides cybersecurity attacks on WANs into two categories: attacks from outside the WAN communications flow and inside attacks that use existing critical WAN paths to target critical assets and systems (see Fig. 2). External attacks typically originate from the same medium through which the network traffic flows. For example, if WAN data are traveling over a wireless medium, an attacker may use another radio to try to monitor or manipulate that link. Internal attacks use the existing critical asset WAN transport mechanisms and communications devices to spread an attack beyond the original exploitation point (either at the centralized control point or the remote outstation). An example of this is an attacker who has compromised the corporate network of a critical asset owner to then use the WAN infrastructure to gain access to remote critical sites.

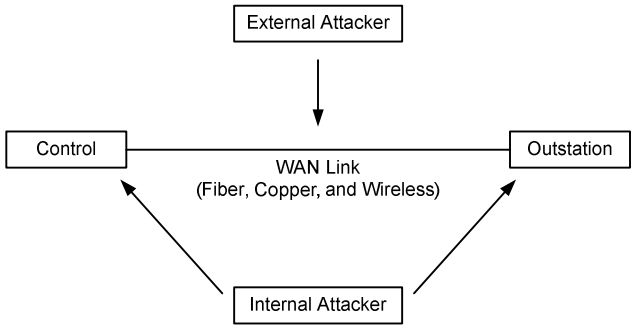

Fig. 2. External and internal WAN attacks.

A. External WAN Attacks

External attacks are further dissected into three attack types: confidentiality, integrity, and availability. Confidentiality attacks encompass all attacks that compromise the secrecy of the data traveling over the WAN. Attackers use well-known sniffing techniques to gather the actual communications data traveling over copper, fiber, and especially wireless links. Not even modern cellular networks are immune to these techniques [2].

Integrity attacks are attempts to manipulate data traveling over the WAN link. A malicious entity can seek to disrupt the authenticity of information by manipulating existing data or injecting new data—even command and control sequences—into the communications stream during transit. Integrity attacks are especially dangerous because they undermine the trust of even the most robust communications networks and possibly cause adverse operations on critical equipment, including trips on protection networks.

Attacks against availability simply seek to disrupt normal WAN services by rendering them unavailable for normal operation. Availability attacks range from the purposeful disruption of traffic flows and malicious cutting of communications cables to mass distributed denial-of-service (DDoS) attacks against utility networks, which is what happened to a German renewable energy operator in late 2012 [3].

B. Internal WAN Attacks

Internal attacks are summarized in this paper as either application-level attacks, which target services and applications that are used for machine-to-machine operations, or user-level attacks, which target services that are used for human-to-machine interactions.

1) Application-Level Attacks

Application-level attacks target services on critical devices and systems that are primarily required for the day-to-day automated needs of critical infrastructure, such as SCADA. Some examples of common machine-to-machine services and protocols include Modbus® TCP and Modbus RTU, DNP3, IEC 61850, the IEEE C37.118 standard for synchrophasors, Microsoft® Remote Procedure Call (MSRPC) for Windows® computers, Microsoft Open Database Connectivity (ODBC), and other proprietary bit- or byte-based serial or Transmission Control Protocol-wrapped (TCP-wrapped) Ethernet protocols. A critical intelligent electronic device (IED) may use one or more of these services at any one time, especially if it uses Ethernet to communicate.

The services used by critical devices and systems can be divided into the following three types:

- Active service is used on a regular basis. It is essential for the short- and medium-term reliable operation of the critical device on which it is used (protection protocols, Modbus, and so on).
- Passive service is used on an irregular basis. It is essential for the medium- or long-term reliable operation of the critical device on which it is used (Simple Network Management Protocol [SNMP]/File Transfer Protocol [FTP] for configuration management and so on).
- Unused services are never used on the critical device, but are left enabled on the device.

When attacking Ethernet networks, malicious entities routinely use automated scanning programs to discover active, passive, or unused services running on critical equipment. Many network scanning tools, such as Nmap (Network Mapper), are freely available for download on the Internet and are often used for reconnaissance before an active cybersecurity attack. Malware exploits security vulnerabilities on existing services to automatically infect large numbers of systems. There are known instances of malware that have actively targeted critical infrastructure by exploiting existing services on Ethernet networks. Stuxnet, perhaps the most visible and effective example, exploited existing services on Windows-based computers and impacted programmable logic controllers (PLCs) at the Natanz nuclear facility in Iran in 2010 [4].

2) User-Level Attacks

User-level attacks are those that attempt to exploit existing user authentication paths to gain or extend access over critical systems and devices. Examples of user access systems range from simple password challenges on critical devices to complex centralized authentication systems that span an entire infrastructure.

User authentication is an access control mechanism that challenges users seeking access to prove their identity. A user can provide proof of identity by providing something the user knows (a password), has (a smart card or cryptographic token), or is (biometric information). The vast majority of authentication systems on critical infrastructure devices involve something the user must know—a password. Password authentication systems are by far the most popular and successful target of attacks. Adversaries attempt to use known default passwords, dictionary word lists, and even brute force combinations of passwords to try to gain access to critical devices or systems. Password reuse also makes this attack more potent, because once a password is discovered, it then can be used to gain access to any device to which that particular user has access.

Authorization rights in access control systems define the limits of access once the user has successfully proven identity (authentication). The goal of attackers is to gain unrestricted access to a system, and therefore, they seek to gain the most privileges possible. If an access control system does not have the capability to restrict the rights of an authenticated user, then the attacker can simply steal the necessary authentication credentials to gain full system access.

User accountability creates an audit trail by recording the user authentication process and, ideally, tracks what an accessing user does while connected to a device or system. This information can then be used for forensics purposes in case an attack or misoperation on a critical system occurs, tracing the problem back to a username. Attackers may try to erase traces or bypass accountability processes to remove the digital trail.

However, attackers can choose to take advantage of the lack of accountability capabilities in critical infrastructure, especially on legacy devices and systems that do not keep information about user access or actions. Many critical infrastructure devices use global shared accounts that do not trace actions back to unique users, which makes true accountability difficult to accomplish without mitigating technologies.

IV. Cybersecurity Best Practice Principles for Wide-Area Networks

There are several effective best practice cybersecurity ideas and techniques that can be used to mitigate the threats mentioned in the previous section, including both external and internal attacks.

A. Use Defense-in-Depth Strategies

It is never a good idea to rely on one method of cyberdefense. By looking at communications networks for critical infrastructure in terms of defense-in-depth strategies, we begin thinking about a layered cybersecurity approach that provides higher resistance to cyberattacks. For an outstation network (whether serial or Ethernet), we discuss four digital defense zones, a physical defense perimeter, and an optional demilitarized zone (DMZ) for special cases. These zones are outlined in Fig. 3.

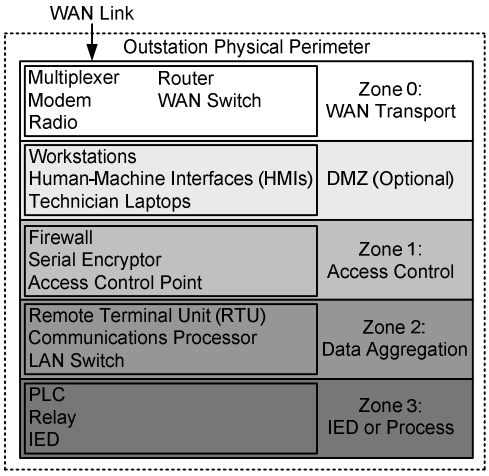

Fig. 3. Apply defense-in-depth strategies to critical networks.

1) Physical Perimeter

The physical perimeter should be a six-walled enclosure (four sides, a roof, and a floor) with some form of access control, whether it is a simple lock and key, a badge, or a biometric user access system.

2) Zone 0: WAN Transport

This zone consists of multiplexers, routers, modems, radios, switches, or other physical devices that provide the communications medium over which WAN traffic flows. Security functions in this zone may include encryption, authentication, and availability, as well as access controls for the WAN device itself.

3) DMZ (Optional)

A DMZ should contain devices or systems that are treated as not trusted. For example, any device that is regularly accessed and controlled by multiple different asset owners should be contained within a DMZ. Devices that are highly vulnerable to hackers should also fall within a DMZ. For example, all Windows- or Linux®-based workstations or computers at critical sites should be in a DMZ. They should never have direct serial or Ethernet access to devices in Zone 3, if technically possible, because consumer-grade operating systems are especially attractive targets for cyberattacks. To compound this issue, it is often a challenge to keep security patches on these operating systems up to date without negatively affecting the reliability or availability of the systems. Therefore, Windows and Linux servers and computers used in critical infrastructure are typically vulnerable to attack by even the least-skilled adversary.

Any DMZ should be included after Zone 0 (wide-area transport) and before Zone 1 (access). However, multiple layers of defense may be built around the DMZ itself, providing the same level of cybersecurity protection as

Zone 1. Therefore, any communication either ingressing or egressing the DMZ is conceptualized as going through Zone 1.

4) Zone 1: Access Control

The access zone filters routable Ethernet communication (Layers 3 and 4 of the Open Systems Interconnection [OSI] model) using technologies such as a firewall. It also authenticates remote devices communicating over a serial link using serial encryptors and provides strong user access control functions, such as authentication, authorization, and accountability, with access proxies.

5) Zone 2: Data Aggregation

The data aggregation zone is where data aggregators, communications processors, and LAN appliances, such as Ethernet switches, are located. This zone may contain its own security functions, including additional password controls and port-level security features on Ethernet devices, such as IEEE 802.1X and virtual LANs (VLANs).

6) Zone 3: IED or Process

This final defense zone includes security controls on the actual critical infrastructure devices themselves. Security functions in this zone are typically not as sophisticated as those in Zones 0 through 2 but may provide additional password protection, alarm contacts, and the ability to disable unused serial or Ethernet ports and services.

The cybersecurity defenses in subsequent zones should adequately answer the question, "What additional security controls are available should the previous zone be compromised or bypassed?" For example, if the physical security perimeter of an outstation is breached and an attacker is able to bypass Zones 0 through 2 by plugging directly into a critical device, which security controls on that device will prevent degradation of reliability or availability?

Two or possibly three of these digital defense zones can be combined in the case of smaller outstations. For instance, an edge router (Zone 0) may also have the functionality of a firewall (Zone 1). However, a good practice is to have the device that performs the duties of Zone 3 be physically separate from the devices that perform the duties of Zones 0, 1, or 2. An example violation of this principle would be directly connecting a critical device (a device that is capable of performing an action necessary for the reliability of critical infrastructure) to the WAN link. Asset operators do not want a PLC or relay—or any critical device that performs or protects a physical process—to be the frontline of a possible security attack, no matter how excellent the security controls local to the device are at that point in time. Exceptions to this rule should only be made for physically protected, dedicated (nonshared) point-to-point WAN links for communications protocols that are not resilient to possible latency added by additional cybersecurity controls. Such a decision must be accompanied by a management-approved cybersecurity risk analysis that proves that the time-sensitivity requirements of the link outweigh the need for cybersecurity protection. In general, the more zones of security defense, the better the overall cyberattack defense. Only Zones 0 and 1 are discussed in detail in this paper.

B. Separate and Filter by Traffic Type

Security controls for communications data flowing into, out of, or across WAN links must be determined by the use and criticality of the data themselves. Engineering access traffic has very different security needs than protection data. Furthermore, protection data have very stringent network latency (both serial and Ethernet) and jitter (Ethernet) requirements that must be met, and certain security controls cannot be applied to protection traffic without negatively affecting the reliability and availability of these data. For example, Ethernet SCADA traffic is easily filtered through a firewall device without negatively affecting the application of the data. However, it is difficult to apply the same security techniques to protection communications data on Ethernet networks without adding costly latency or jitter. Time-division multiplexing (TDM), when combined with multiplexer devices, is a solution discussed in depth in Section V that allows multiple communications streams to be secured and sent over a WAN, regardless of communications data type.

C. Encrypt and Authenticate Data on the WAN Link

One of the best methods to mitigate external WAN attacks is to use encryption and authentication of data on the WAN link between the enterprise and all outstations or between outstations themselves. Encryption, the process by which data are scrambled so adversaries cannot analyze the data flowing on the WAN, helps mitigate confidentiality attacks. Authentication, which helps prevent injection or manipulation of data, is an excellent mitigation of integrity attacks.

D. Integrate Strong Controls for User Access

Strong authentication, authorization, and accountability are essential for any remote or local engineering access system, both for WAN transport devices themselves and critical infrastructure devices. Robust access controls are necessary to prevent or mitigate password attacks, restrict the privileges of accessing users, and provide an audit trail that can be used effectively in case of system misoperation or cyberattack. Strong access controls can also be used for legacy critical systems and devices, as long as all access to those systems and devices is funneled through a proxy device that can add the necessary controls before a user is allowed to access the critical device.

V. Securing the External WAN

A. Segregating Application Traffic Provides Increased Security for Mission-Critical Services

Two modern transport technologies used for WAN communication are TDM or packet-based communication [5]. Both technologies have features for securing the data being transported across the network.

TDM divides the transport bandwidth into a series of time slots, each with a specific payload size. An analogy for TDM is first-class reserved seating on a regular high-speed train. Each seat on the train is reserved for a different application or service, and only data from that application are allowed to occupy those seats. Each application is allocated the same reserved seat on every train. Because the seat allocation is known for each application, the application receiving data from the train does not need to look at the data contents of every seat to determine if the application is the intended recipient. The routing and connectivity of data through a TDM system are managed by a dedicated operations, administration, maintenance, and provisioning (OAM&P) layer that manages the seating reservations for each application.

Ethernet is one of the most widely implemented packet-based transport technologies. Unlike TDM, Ethernet does not use the concept of preallocated seats or timeslots to send data. Instead, all applications share the same bandwidth. Ethernet operates like a road system with trucks and cars carrying the data. The traffic is "bursty," payload sizes vary, and access to the shared bandwidth is random.

The WAN is a shared medium for transporting large volumes of data between locations. Some WANs are operated by companies that lease bandwidth to other users. Data segregation is an important attribute in network management because it controls how bandwidth is allocated to different users and limits which network devices or ports have access to specific data. Data segregation is also important for network security. If the same physical WAN is shared between different organizations, it is essential that data traffic from one company cannot be accessed by another company on the same network. Similarly, in networks carrying data for critical systems, it is important to segregate protection traffic from noncritical IT traffic. Data can be physically segregated using different fibers or electrical wires for different services or logically segregated using protocols within the shared fiber.

TDM provides security by segregating data into separate timeslots and transporting the data to dedicated end points or ports. The end user or application only sees data that are intended for its use.

In Ethernet, all data are transmitted over shared bandwidth. Ethernet uses security methods for encrypting the contents of an Ethernet packet or frame to prevent an attacker or an unintended recipient from reading the contents of each Ethernet frame. Virtual private networks (VPNs) provide secure point-to-point connections through public and private networks by encapsulating and encrypting the entire Ethernet frame within an outer VPN packet. Ethernet supports the segregation of traffic through the use of VLANs. With VLANs, it is possible to partition and direct data at OSI Layers 2 through 7 to specific ports or network devices. Because of their traffic separation properties, VLANs are effective at mitigating the spread of cyberattacks that use flat, nonrouted networks. An adversary may gain access to and negatively affect devices and services on one VLAN but be prevented from gaining access to the system as a whole without breaking into separate defenses for other VLANs.

It is possible to combine the attributes of TDM and Ethernet by running Ethernet over TDM. This approach combines the flexibility of Ethernet with the dedicated network management and traffic segregation characteristics of TDM.

Synchronous optical network (SONET) is a TDM standard that is widely used in the United States and Canada. Synchronous digital hierarchy (SDH) was defined by the European Telecommunications Standards Institute (ETSI) and has been adopted by the rest of the world. Both standards support the ability to run Ethernet over TDM.

Fig. 4 shows how segregated data pipes can have allocated, isolated bandwidth for specific services or applications within the same physical fiber using the SONET transport structure. These pipes can be Ethernet or native TDM.

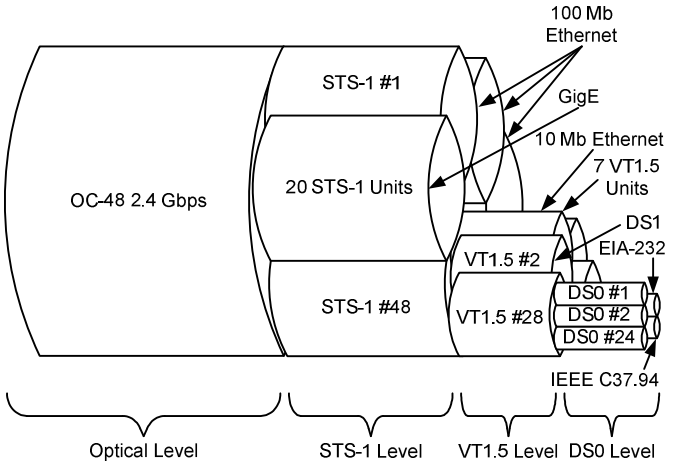

Fig. 4. Mapping TDM and Ethernet pipes into the SONET transport structure.

A DS1 can be mapped directly into a VT1.5 (where VT is virtual tributary). A 10 Mb Ethernet signal can be mapped into 7 VT1.5 units. A 100 Mb Ethernet signal can be mapped into 2 STS-1 units (where STS is synchronous transport signal), and a GigE (gigabit Ethernet) signal can be mapped into 20 STS-1 units.

Table I, which shows the SONET hierarchy, and Fig. 4 illustrate how different bandwidth Ethernet channels can be mapped into the SONET framing structure.

TABLE I
SONET DIGITAL HIERARCHY

Level	Line Rate (Mbps)	Number of 64 kbps Channels	Number of DS1 Units
VT1.5	1.728	24	1
STS-1	51.84	672	28
OC-1	51.84	672	28
OC-3	155.52	2,016	84
OC-12	622.08	8,064	336
OC-48	2,488.32	32,256	1,344
OC-192	9,953.28	129,024	5,376

The data in any Ethernet or TDM pipe are segregated from the data in any other pipe. In the case of a leased WAN service, each pipe would be dedicated to a different user or organization. In critical infrastructure applications, each pipe would be dedicated to specific applications, such as SCADA, relay protection traffic, or IT services, as shown in Fig. 5.

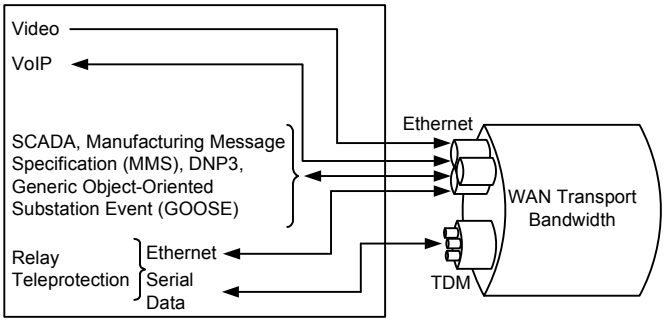

Fig. 5. Segregation of application traffic using TDM and Ethernet pipes.

The ability to segregate data within the same fiber to specific services that map to dedicated end devices or ports provides greater security for the system. The OAM&P function provides the network manager with event and alarm information for each data pipe and physical port provisioned on the network. For example, the network manager is made aware if the technician responsible for the video camera network pulled a cable from a port that was not part of the video network. This could be the result of an innocent mistake or a deliberate attempt to cause harm, but the instant notification to the network manager allows the situation to be detected and appropriate action taken. The other advantage of isolating bandwidth into segregated pipes is that it prevents high-bandwidth Ethernet applications from consuming all the available resources and reducing network throughput for other services. This approach can reduce the impact to a network during a denial-of-service (DoS) attack. In a DoS attack, the perpetrator floods the bandwidth of a target device by sending large numbers of Internet Protocol (IP) packets, which slows or crashes the network. Running Ethernet over TDM and segregating traffic into isolated pipes restricts the DoS attack to the services running on the target Ethernet pipe and the other local Ethernet services running on the attacked Ethernet switch. The DoS attack saturates the Ethernet switch message buffers in the target node, affecting the target Ethernet pipe and all local Ethernet services, but the attack does not saturate Ethernet traffic traveling in other Ethernet pipes that are running on the SONET line transport between other nodes on the network. In addition, TDM allows the network manager to control which Ethernet pipes are dropped at each node, allowing critical services to only be dropped at selective nodes. This restricts the number of Ethernet pipes that an attacker can target from a single node.

B. WAN Line Encryption

A man-in-the-middle attack is a specific form of cyberattack where the attacker is able to intercept network traffic without detection for the purpose of eavesdropping (confidentiality attack) or inserting or modifying messages (integrity attack). Fig. 6 shows an attacker gaining access to the nonsecured WAN for the purpose of performing a man-in-the-middle attack.

The solution to this type of threat is to encrypt and authenticate all data being sent across the WAN transport network to prevent the hacker from interpreting application data and injecting rogue messages.

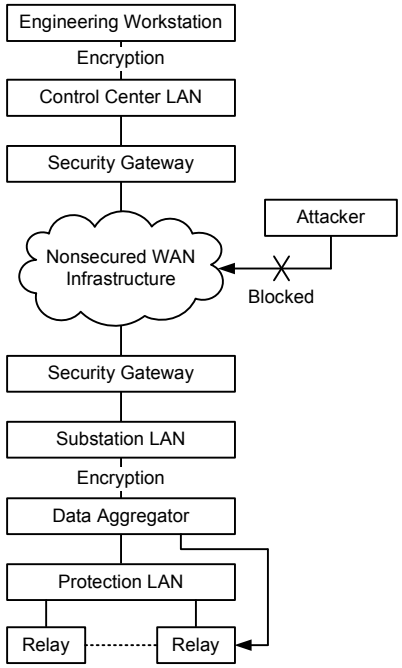

Fig. 6. Man-in-the-middle attack.

The following are three broad approaches used to secure data at the WAN level:

- Perform encryption at the end application.
- Use a VPN to create a secure tunnel through the WAN transport network.
- Perform encryption at the transport level, and encrypt all the data in the datagram.

Best practice recommendations are a combination of all three approaches to ensure the highest level of security. Organizations that have traffic that travels across external networks that are run by other companies should adopt end application encryption and VPNs to secure data. All external networks should be treated as nontrusted because the implementation of security best practices cannot be guaranteed. To protect the network infrastructure of an organization against the threat of a man-in-the-middle attack, the recommended approach is to encrypt at the WAN transport level. With end application encryption, it is still possible for the hacker to gain some knowledge of the network. The IP addresses from Ethernet header information can still be read, giving clues concerning where data are traveling to and from. Looking at data patterns based on frequency, time of day, and packet sizes can give more specific information concerning the types of devices that are communicating, allowing a hacker to time a DoS attack when the most critical data are being transported. Encrypting at the WAN transport level prevents the attacker from gaining any information on the

network traffic at all. In the past, WAN line encryption required dedicated hardware that was expensive and typically limited to core network carrier-class equipment. This is no longer the case, and cost-effective line encryption is now available for ruggedized substation WAN multiplexer equipment on the edge of the network. For most power utility organizations, the core network WAN provides communication among corporate offices, data centers, and network operation control centers. The edge network provides communication to generation plants, substations, and transmission and distribution infrastructure. These network edge facilities are significant targets for cyberterrorists. Implementing WAN line encryption on the edge network provides the greatest protection against man-in-the-middle network intrusion.

C. Securing Network Management System Interfaces

All WANs require a network manager function to configure and monitor the communications equipment and manage the operation of the network. The network manager is typically a software program running on a computer or workstation that connects into the WAN via a network node using a TCP/IP communications interface. The network management system (NMS) is a critical component of the network infrastructure, and it enables the authorized user to change any parameter or setting on the network. For that reason, all user access to the NMS must be controlled via user authentication, authorization, and accountability procedures. In addition, the communications interface that the NMS uses to connect to the network must have safeguards that are resistant to unauthorized users attacking the management interface.

SNMP is an industry standard protocol for managing network devices. The latest version at the time of publication, SNMPv3, supports improved authentication and cryptographic security to secure the data exchanges between the network management system and the network equipment.

SNMPv3 provides the following security capabilities:

- Encryption of packets to prevent sniffing or eavesdropping.
- Data integrity to ensure a data packet has not been interfered with.
- Authentication that a message is from a trusted source.

VI. Securing the Internal WAN

Once adequate cybersecurity controls have been applied to mitigate external WAN threats, the controls needed to thwart application-level attacks and user-level attacks as a result of internal threats need to be understood.

It is dangerous to assume that if the WAN link incorporates security features such as encryption, authentication, and availability, the majority of cyberthreats are mitigated. To thwart application- and user-level attacks, critical infrastructure asset owners should use strong user access controls and traffic filters at each physical outstation, no matter how the outstation is connected back to a central command point.

The reasoning for this is simple: if the central command-and-control point for the critical infrastructure network is compromised or any outstation network is compromised, a system without layered security controls is open to an attack spreading over internal WAN links. If outstations are connected back to the central control point using a star architecture and no additional security controls exist at the outstation, an internal attacker may effectively compromise and attack all outstations by successfully attacking the central control point (see Fig. 7).

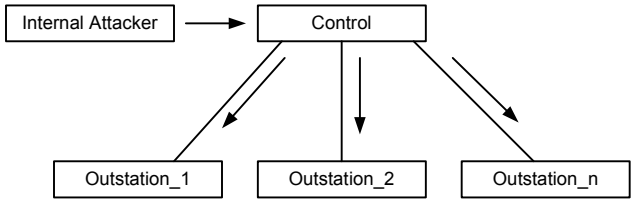

Fig. 7. Attacker compromises outstations via WAN from central control point (star architecture).

Furthermore, it is possible for the attacker to compromise the critical network at a single outstation in order to gain access to the control point and possibly jump to other control points, as shown in Fig. 8.

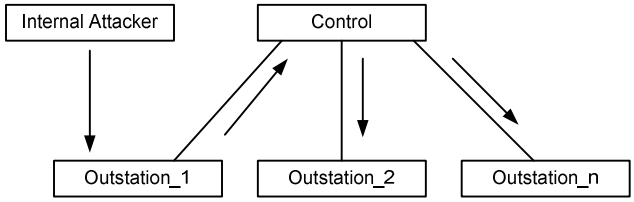

Fig. 8. Attacker compromises control and outstations via WAN from outstation (star architecture).

If outstations are connected back to the control point via a ring architecture and there are inadequate security controls in place, an adversary can simply attack any outstation in order to gain access to the ring network. From there, an attacker can possibly compromise any other outstation connected to that ring—or even the control point itself (see Fig. 9).

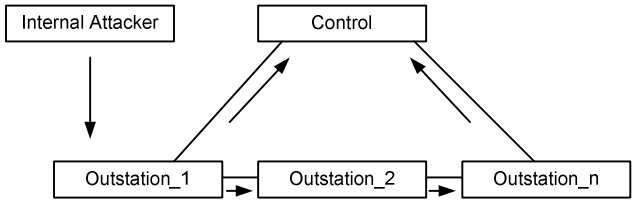

Fig. 9. Attacker compromises control and outstations via WAN from outstation (ring architecture).

Knowing where to place security controls is not a difficult exercise when using cybersecurity best practices. Traffic separation and filtration, strong access controls, and defense-in-depth strategies are all common techniques that have been vetted by IT experts for years and, when used effectively, are good techniques for thwarting cyberattackers. An example of a typical nonsecure Ethernet-based outstation network is shown in Fig. 10.

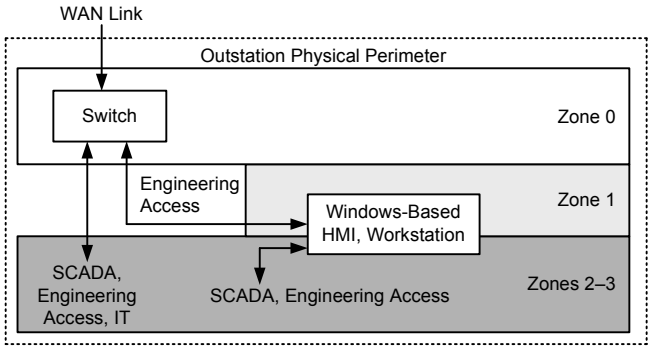

Fig. 10. Example of a nonsecure Ethernet-connected outstation network.

When comparing this design with possible application- or user-level attacks, the network has the following vulnerabilities:

- There is no filtering of incoming or outgoing Ethernet SCADA or IT data. Even if the WAN link itself is encrypted, an attacker can compromise the outstation network or the central control network and initiate service-level attacks on SCADA and IT devices. If passive or unused services are enabled on critical devices in Zones 2 or 3, an adversary can attempt a DoS attack on those services in an attempt to affect the availability of those devices or cause a misoperation.
- A Windows-based computer at the critical site doubles as both an HMI and workstation for engineering access. The computer, which has not been patched in several months, is directly connected to critical IEDs. Malware coming from an infected device at a central control point or piggybacking on a USB drive used by a technician can attack and compromise the Windows machine. Because the computer is directly connected to Zone 3, the critical IEDs placed there are directly accessible to the attack.
- Engineering access is not funneled through the Windows-based computer, where proper access controls may be used. Instead, engineering access to Zone 3 itself is available from the same direct network as the SCADA and IT devices. The more robust user access controls available on the Windows machine can be bypassed, and the security of user access is entirely reliant on the critical devices themselves.

The cybersecurity of the outstation is strongly enhanced by redesigning this network using solid defense-in-depth practices and forcing all user access through a proxy device, as shown in Fig. 11.

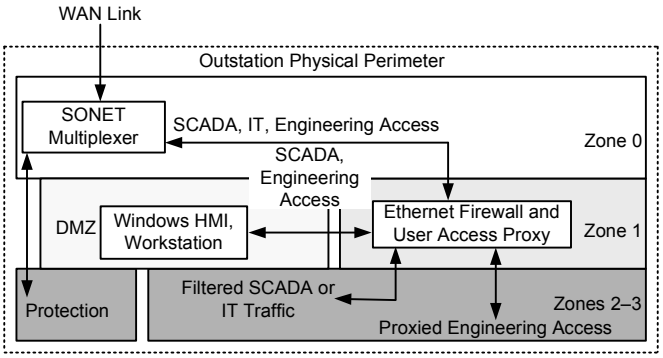

Fig. 11. Ethernet-connected outstation network with cybersecurity best practices.

While this example uses a SONET multiplexing device on the WAN for the most reliable protection traffic, the same idea can be achieved for engineering access, SCADA, and IT traffic using an Ethernet-based WAN device. The following lists the benefits of this design:

- The encrypting SONET multiplexer separates the different traffic types into different TDM pipes. Protection traffic is either physically separated (if using serial) or logically separated (if using Ethernet) by using VLANs. This way, serial traffic neither affects or is affected by the SCADA, engineering access, or IT communication running over separate Ethernet-tunneled TDM. VLAN-separated Ethernet traffic can still influence the availability of other traffic on the same switch, but these effects are minimized by using proper port speeds, rate limiting, Ethernet traffic prioritization, and Layer 2 filtering techniques.
- All SCADA, IT, and engineering access data are filtered through an Ethernet firewall in order to minimize service-level attacks on devices in Zones 2 and 3.
- All engineering access to devices in Zone 3 is filtered through a user access proxy. The proxy enforces additional strong authentication, authorization, and accountability to critical IEDs that cannot offer the same robust cybersecurity controls.
- The Windows workstation is confined to a DMZ and treated as an untrusted device. All user access and SCADA from the workstation must be filtered through the firewall or proxy in order to access Zones 2 or 3.

Now, even if the central control or outstation network is compromised, it is much easier to isolate that threat to the local exploitation point and resist the ability of the threat to spread to other zones or sites.

The same cybersecurity principles can be applied to legacy serial-based, dial-up infrastructure. By simply applying user access control and defense-in-depth ideas, we can mitigate threats against dial-up critical infrastructure.

Fig. 12 shows how these security principles can be used for serial SCADA or serial protection communication. In those instances, the serial encryptor is providing authentication of the remote SCADA master or other IED, rather than of the remote user.

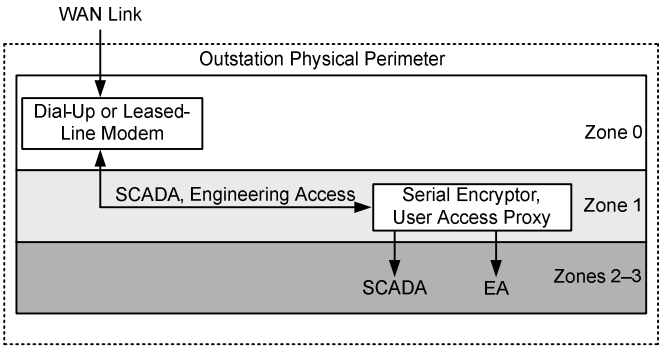

Fig. 12. Dial-up accessible outstation network with cybersecurity best practices.

The dial-up modem connects to a serial encrypting transceiver instead of directly to Zones 2 or 3. The serial encrypting transceiver is used to authenticate the user who is dialing in and provide an audit trail. Additional user access controls may be applied at Zone 1. For instance, the serial encryptor can direct the authenticated user to another user access proxy, through which additional cybersecurity controls may be applied, such as authorization rights.

VII. System Monitoring

Given enough time, money, and expertise, a sophisticated cyberadversary can find a way to attack even the most secure network or device. Network monitoring provides an additional layer of defense and the ability to detect hardware failures, security breakdowns, and active attacks. Monitoring involves actively analyzing traffic on the network to look for suspicious system activity, such as unusual traffic, unauthorized user activity, and unusual device configuration changes. The Syslog protocol has become the standard logging solution for network monitoring and can be found on many network systems. Syslog implementations exist for a variety of different operating systems and are commonly found in network devices, such as switches, firewalls, and WAN communications equipment. Syslog creates a separation between the software subsystems that generate, store, report, and analyze messages. It can be used for network system management and security auditing, as well as providing a tool for network analysis and debugging. Syslog can be used to integrate log data from many different types of systems into a central repository.

Syslog can also be used to provide an audit trail by generating security-relevant chronological records that contain documented evidence of the sequence of activities affecting any specific operation, procedure, or event. Audit records with accurate time stamps of individual users are extremely useful for identifying compromised accounts and reversing the damage done by cybersecurity attacks. Having an accurate time reference for time-stamping logged events on the network is extremely important in the analysis of syslog data. Many networks make use of high-accuracy Global Positioning System-synchronized (GPS-synchronized) clocks or cesium or rubidium references to provide network synchronization. These clock references provide the basis for a time-distribution solution for syslog message logging. Many time-distribution protocols can be used to provide time-stamping with sufficient accuracy for message logging. The more commonly used protocols are Network Time Protocol (NTP), Precision Time Protocol (PTP), and IRIG-B.

Logging and alerting processes on critical devices and networks should supervise and record the actions of all users during an active session. Security information and event management (SIEM) solutions offer real-time security analysis performed by hardware or software algorithms to identify potential threats. SIEM solutions are capable of consolidating and correlating recorded events into useful charts or diagrams to enable system managers to quickly identify unusual network traffic patterns that fall outside a standard profile.

VIII. Conclusion

There have been successful cyberattacks against critical infrastructure both in and outside of North America. Government and private organizations are establishing programs to put policies and procedures in place to protect critical infrastructure cyberassets from the threat of cyberattacks. Looking forward, WAN technology will be increasingly used for the control and automation of utility and industrial processes, with the electric power system being one of the largest adopters. With the growth in the usage of WAN communications technology, the need for effective cybersecurity measures is becoming increasingly important. The best practice cybersecurity approaches discussed in this paper are manageable and implementable using technology that is commercially available today. Used together, the methods outlined provide a practical approach to mitigate the risks of cyberattacks.

IX. References

[1] D. Dolezilek, "Case Study of Mission-Critical Smart Grid Remedial Action Schemes Via Ethernet," proceedings of the 37th Annual Western Protective Relay Conference, Spokane, WA, October 2010.

[2] K. J. O'Brien, "Cellphone Encryption Code Is Divulged," *The New York Times*, December 28, 2009. Available: http://www.nytimes.com/2009/12/29/technology/29hack.html.

[3] A. Nelsen, "European Renewable Power Grid Rocked by Cyber-Attack," *EurActiv.com*, December 20, 2012. Available: http://www.euractiv.com/energy/european-renewable-power-grid-ro-news-516541.

[4] "A Worm in the Centrifuge," *The Economist*, September 30, 2010. Available: http://www.economist.com/node/17147818.

[5] E. O. Schweitzer, III, D. Whitehead, K. Fodero, and P. Robertson, "Merging SONET and Ethernet Communications for Power System Applications," proceedings of the 38th Annual Western Protective Relay Conference, Spokane, WA, October 2011.

X. BIOGRAPHIES

Paul Robertson is a lead product manager for the wide-area networking group at Schweitzer Engineering Laboratories, Inc. (SEL). He has over 20 years of experience developing and marketing products for the telecommunications industry, spanning cellular wireless and wire line communications systems. Paul worked in various technical and marketing roles for Motorola, Hewlett-Packard, and Agilent Technologies before joining SEL. He has a BEng in electrical and electronic engineering from Strathclyde University and an MBA from Edinburgh Business School.

Colin Gordon is an application engineer within the research and development group at Schweitzer Engineering Laboratories, Inc. (SEL), specializing in cybersecurity solutions and services for critical infrastructure. His work experience includes secure network design, implementation, testing, and North American Energy Reliability Corporation Critical Infrastructure Protection (NERC CIP) compliance consultation for utilities and asset owners in North America and abroad. Colin joined SEL in January 2008 as a product management intern and holds a bachelors degree in computer engineering from the University of Idaho.

Simon Loo is an integration application engineer at Schweitzer Engineering Laboratories, Inc. (SEL). Simon received his B.S. in electrical engineering from the University of Arizona with a minor in math and computer engineering. He joined SEL in 2011 as an associate integration application engineer. Prior to joining SEL, he was with Southwest Transmission Cooperative, Inc. as a field engineer and Freeport McMoRan Cooper & Gold, Inc. as an electrical engineer. He is a member of IEEE and was vice-chair of his IEEE student branch in 2008.