Synchrophasor Security Practices

John Stewart, *Tennessee Valley Authority*
Thomas Maufer, *Mu Dynamics, Inc.*
Rhett Smith, Chris Anderson, and Eren Ersonmez, *Schweitzer Engineering Laboratories, Inc.*

***Abstract*—Synchrophasors provide a powerful tool to change the view of the power system from an estimated state to a directly measured state. Synchrophasor solutions are wide-area systems that may include many phasor measurement units (PMUs), phasor data concentrators (PDCs), and visualization, archiving, and control systems. These system components continuously communicate across large geographical areas, which may involve communications links over untrusted channels. In some cases, synchrophasor systems intersect with protection systems; such is the case when PMU functionality is integrated in protective relays. Therefore, cybersecurity risks must be addressed to ensure that reliable operations are maintained. A common concern is that enabling synchrophasors may introduce cybersecurity risks to protection systems. This paper examines the potential risks arising from using synchrophasor data communication from a substation and provides suggestions on how to mitigate the risk of cyberattacks while preserving the benefits of synchrophasors. In addition, this paper discusses how to approach risk assessments when enabling synchrophasors in substations and identify solutions for access control, confidentiality, data integrity, and availability. Finally, this paper provides test results showing data performance and security robustness with and without the security safeguards applied.**

I. Introduction

Synchrophasor technology can be summarized as marking power system quantities with a high-accuracy time tag in order to be able to use data from multiple sources in a coherent manner [1].

Traditional supervisory control and data acquisition (SCADA) systems can gather measurements at intervals ranging between 2 and 15 seconds. This level of resolution, plus the fact that data are sampled and received asynchronously, requires that the information be processed by very complex algorithms before the power system state can be determined. However, as a result of the methods of data collection, the calculated power system state is an approximation because there are inherent errors due to possible system changes during the SCADA scan times. Hence, power system state estimators are used to periodically calculate the power system state, which typically includes bus voltage magnitude and angle at various locations. These quantities are then used to make decisions about how to operate the electric power system.

In contrast, synchronized phasor measurements are typically collected at 30 measurements per second. Newer systems that include control applications may have measurement rates between 60 and 240 measurements per second. Furthermore, synchrophasors use a high-precision common time source that provides 100-nanosecond accuracy at each measurement location. Today, most synchrophasor systems use Global Positioning System (GPS) signals as this common time source. This accuracy allows the measurement of power system quantities at different locations across a wide-area system at the same instant in time. Because these data are measured contemporaneously with a common time reference, the measurements can be compared directly, without the need for complex algorithms. The system state is measured, allowing for better-informed decisions regarding how to operate the electric power system.

Electronic communication is an important component of almost all synchrophasor implementations and brings associated concerns about cybersecurity. Common concerns include new communications links opening the way for cyberattacks on existing systems or on the new synchrophasor system itself. Because most of the power system is monitored and controlled electronically, a successful cyberattack could cause great physical and financial damage. Therefore, it is essential that cybersecurity risks be carefully analyzed and mitigated.

This paper presents some of the best practices related to mitigating cybersecurity risks in synchrophasor systems, based on experience with implementations at utilities such as Tennessee Valley Authority (TVA) and other organizations. These practices are not one-size-fits-all solutions for all applications. However, one or more of these practices can be used if the methods are technically and economically feasible for the application.

II. Background

A. Synchrophasor System Components

A device that measures electrical quantities from the power system is called a phasor measurement unit (PMU). PMU functionality can be included in protective relays, meters, digital fault recorders (DFRs), or other intelligent electronic devices (IEDs). In fact, PMU functionality has been included in IEDs since the early 2000s; therefore, tens of thousands of IEDs that include PMU functionality are already in service in the power system. These PMUs need to be connected to GPS-synchronized clocks that provide a high-precision common time source. Synchrophasor data can then be sent to end-user applications for archiving, monitoring, or control.

Applications can receive data from each individual PMU or in a concentrated format from a phasor data concentrator (PDC). PDCs receive data from multiple PMUs and time-align the data based on the measurement time tags. For each time tag, the PDC generates a concentrated packet (also known as a "super packet") that includes the values from all PMUs. Then,

the PDC serves these data to other devices or applications, typically using network communication. PDCs may include other functionality in addition to data concentration. For example, a PDC can perform calculations and logical operations, archive data, or even take control actions.

Fig. 1 shows typical synchrophasor system components and data flow.

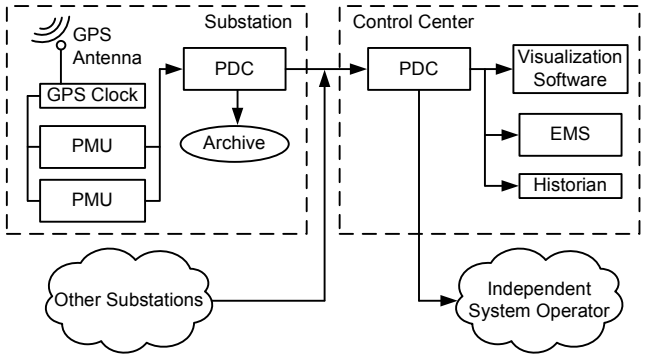

Fig. 1. Example synchrophasor system

B. Synchrophasor Applications

Although the first synchrophasor systems were introduced in the early 1990s [2], this is still a new technology to many utilities. Especially after the 2003 blackout in the eastern United States and Canada, it has become apparent that the power system is growing more complex every day. We need better ways to manage the grid in order to quickly respond to problems or prevent them from occurring in the first place. Lately, more utilities are looking at synchrophasors as part of the solution. There are a rapidly increasing number of utilities that are experimenting with synchrophasor technology in their test labs or pilot field projects. Several utilities are already using synchrophasors in day-to-day power system operations and planning. Common usages include:

- Wide-area visualization
- Modal analysis
- Post-event analysis
- Power system model validation
- Wide-area synchronism check
- Loop flow analysis

There are also a small, but increasing, number of applications in test or operation where synchrophasors are used in automated real-time control systems called special integrity protection schemes (SIPS). For example, in one SIPS implementation, the synchrophasor system identifies undamped oscillations and automatically takes corrective action before the system collapses [3]. Another example is the Southern California Edison synchrophasor system that automatically controls their static VAR compensator (SVC) system for maintaining voltage stability [4].

C. Synchrophasor Communication

Today, the most dominant protocol for communicating synchrophasor data is IEEE C37.118, which includes requirements for synchrophasor measurements as well as a data communications protocol for exchanging synchrophasor data in real time [5]. Less common protocols include PDCStream, Fast Message, and IEC 61850.

There is currently an IEEE and IEC joint team working on synchrophasor mapping for IEC 61850. There are also ongoing efforts to create an Internet Protocol (IP) profile for IEC 61850, which will allow IEC 61850 communication in wide-area systems. When these efforts are completed, we may start seeing more cases of IEC 61850 for synchrophasors.

IEEE C37.118 defines a binary messaging format, but it does not require any specific medium or transport mechanism for communicating these frames. Most implementations today use Transmission Control Protocol/Internet Protocol (TCP/IP), User Datagram Protocol/Internet Protocol (UDP/IP), or EIA-232. EIA-232 is most commonly used within the substation or for transporting a small set of data from a single PMU between substations and control centers. In a wide-area system where there are multiple PMUs involved, IP networks are generally employed. Substations may be linked to each other and control centers by leased lines, privately owned synchronous optical networks (SONETs), wireless links, and so on.

There are four types of IEEE C37.118 message frames:

- Data
- Configuration
- Command
- Header

Clients send command frames to either request configuration and header frames or to start or stop the data stream. Data frames include the actual data being sent in a raw binary form in an integer, floating-point, or Boolean format. Configuration frames define the data frames so that clients can know how to interpret the raw bytes. Header frames are not commonly used today, but they are intended to transmit any general information about the PMU or the PDC in a text form. IEEE C37.118 frames include common header fields (not to be confused with the header frame) in the beginning and a cyclic redundancy check 16 (CRC 16) checksum at the end. Header fields include the PMU/PDC ID and a time tag. A typical IEEE C37.118 conversation is shown in Fig. 2. The client requests a configuration frame from the PMU. After receiving the configuration frame, the client asks for the data stream. The PMU continues to send data frames until the client requests that it stop the data stream.

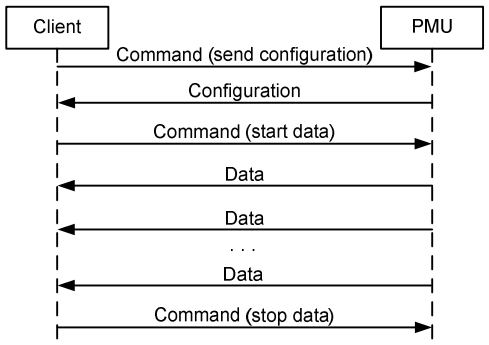

Fig. 2. Example IEEE C37.118 conversation

D. Importance of Security

For many organizations, synchrophasor systems are, or will be, critical for power system operations. In some other organizations, synchrophasor systems are not necessarily considered critical today, because they are either in experimental stages or only complement the existing systems. Although synchrophasor systems are not critical in all organizations, we do see an upward trend in the use of synchrophasors in protection, operations, and planning. In the future, it seems likely that the role of synchrophasor systems will become critical for all utilities and independent system operators. However, these critical systems probably will not be built from scratch. Instead, there will be gradual transitions from existing systems. Most new implementation projects will try to use the existing systems and infrastructures to their maximum extent for time- and cost-efficiency. Therefore, it is probable that the noncritical systems of today will become a part of the critical systems of tomorrow.

If security is not built into the existing systems from the ground up, there is the risk of leaving security vulnerabilities in place when implementing new critical systems. For this reason, even if the synchrophasor systems being implemented today are not considered critical for power system operations, it is advisable to treat them as critical systems from the beginning when designing security. Although this does not mean that every system needs to be classified as a critical cyberasset in terms of compliance with government security requirements, such as Critical Infrastructure Protection (CIP) compliance, it is a good business practice to apply as many of the CIP requirements as are technically and economically feasible. Doing so not only reduces the risk of having security vulnerabilities in critical systems in the future but also makes compliance to these requirements easier should this be required for these systems in the future.

III. Concerns and Solutions

In order to analyze and mitigate cybersecurity risks related to synchrophasor systems, we can group the concerns into the following two categories:

- Substation security
- Information security

A. Substation Security

An obvious concern is that an attacker will gain access to the substation network and the cyberassets within it. A complete security assessment for a substation should include physical security mechanisms and other procedures against an attack within the physical perimeters of the substation. These include video surveillance, physical access control, perimeter fencing or walls, personnel training, background checks, and so on. Physical attacks or cyberattacks within the substation are outside the scope of this paper. In this paper, we focus on attacks from external cyberspace.

To protect the substation network from external attacks, a general best practice is to form an electronic security perimeter around the substation network. Access points into this perimeter are continuously and closely controlled, because these access points are where the attacks will come from. Therefore, a best practice is to minimize the number of these access points and also minimize their exposure to the outside world.

In order to reduce the likelihood of an attack directly on the PMU, it is a good security practice to have the PMUs multiple layers deep in the substation network architecture, each layer providing an additional level of security. In a multilayered architecture, the synchrophasor data that the PMUs send out of the substation pass through multiple protective layers before leaving the substation. At the same time, accessing the PMU from the outside requires passing through these same multiple layers. This makes it much harder for an attacker to gain access to the PMU (or other IEDs) within the same local-area network (LAN) of the PMU.

Fig. 3 shows a multilayered architecture. The access point through which synchrophasor data are sent outside the substation resides in the first layer: the security gateway device. The security gateway fulfills two complementary roles:

- Firewall
- Virtual private network (VPN) tunneling

We will talk about VPN tunneling more in Section III, Subsection B.

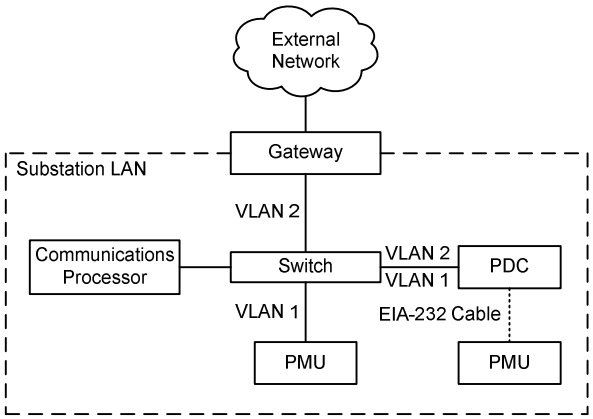

Fig. 3. Multilayered architecture

Firewalls restrict the incoming and outgoing network traffic based on a set of user-defined rules or policies. When setting up the firewall rules, we highly recommend using a white-list approach (also known as deny-by-default). In a white-list approach, all traffic is blocked unless explicitly allowed by a rule. This is to the opposite of a black-list approach, where all traffic is allowed unless explicitly blocked by a rule.

Using the white-list approach, the only traffic allowed should be that which is necessary to deliver the synchrophasor data to the clients outside the substation. Devices behind the security gateway should be invisible to entities in the external network, except those known clients.

In this multilayered architecture, PDCs also act as an additional layer of security. Synchrophasor data clients outside the substation receive data from a PDC rather than directly from PMUs. This single output point makes it simpler to secure the network traffic between the substation and external

clients. Firewall rules can limit the synchrophasor data and settings traffic to a single device instead of multiple PMUs.

Deploying a PDC at the substation also minimizes the need for outsiders to access PMU settings. For example, if the IP address of the control center PDC changes, it is sufficient to remotely access the PDC settings rather than having to access the settings of each individual PMU. In this way, it is possible to give a group of people access to edit synchrophasor output settings without giving any of them access to the PMU settings.

Additionally, having a multilayered security architecture makes it possible to enhance the synchrophasor security mechanisms with a single point of upgrade if new security features are available. For example, IEEE C37.118-2005 is the most common protocol for communicating synchrophasor data today, yet there are almost no inherent security features in this protocol because it leaves security implementation to other network layers. In the future, if there were enhancements to IEEE C37.118 for implementing new security features (such as public key infrastructure [PKI] and role-based access control [RBAC]), it would be much simpler to implement these enhancements through an update to the PDC, rather than having to upgrade PMUs.

Similarly, this multilayered architecture allows the implementation of cutting-edge security mechanisms, even if PMUs do not directly support them. Security mechanisms can be applied through wrapper protocols such as Internet Protocol Security (IPSec). If there were any enhancements to the network traffic filtering or encryption, such enhancements could be implemented at the security gateway very efficiently, without disturbing the PMUs. On the other hand, if these features were implemented directly at the PMU level, we would have to keep all PMU devices updated with the ever-changing security technology. As we can see, numerous operational benefits accrue from a multilayered architecture.

Finally, the PMUs reside behind the PDC. The PMUs can be connected to the PDC via EIA-232 or Ethernet.

If there are several devices, such as human-machine interfaces (HMIs), communications processors, and substation PCs, that require remote access from different groups in the organization, we may choose to separate these devices into a perimeter network, as shown in Fig. 4.

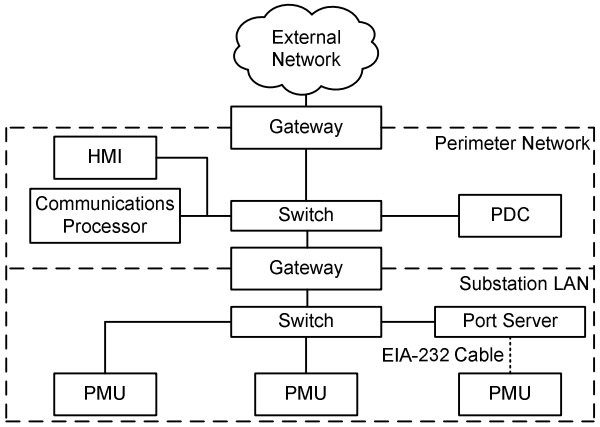

Fig. 4. Multilayered architecture with a perimeter network

In the architecture shown in Fig. 4, we establish a separate electronic security perimeter for devices directly connected to the power system, which gives an additional level of security. Devices that serve external clients reside within a perimeter network. A security gateway device separates the substation LAN and the perimeter network. It is important that all communications route through this local gateway and that there is no bridging or other connection between the two networks. For example, if a communications processor with multiple Ethernet ports is connected to the perimeter network on one port and the substation LAN on the other port, this is a dual-homed host and potentially allows someone who has access to the communications processor to bypass the security gateway. Similarly, an EIA-232 connection between a PMU and a device in the perimeter network bypasses the access controls in the gateway. If the PMU does not have an Ethernet interface, we recommend using a port server or an Ethernet transceiver to convert EIA-232 traffic to Ethernet so that it can be routed through the security gateway, as shown in Fig. 4.

1) Unidirectional Synchrophasor Streams

Especially for cases where synchrophasor data traffic flows over untrusted networks, we propose using a unidirectional streaming mechanism using IEEE C37.118, which we call UDP Secure (UDP_S).

The following points explain the UDP_S protocol behavior:

- The server (data sender) sends IEEE C37.118 data frames in UDP datagrams to the client (data receiver).
- The server neither expects nor accepts any incoming data from the client. Therefore, IEEE C37.118 command frames are not used in this protocol (they are disabled).
- The server publishes its IEEE C37.118 configuration when it is activated, when the configuration changes, and at the top of each minute.
- The server is agnostic of client states. As soon as it is activated, it tries to send data and configuration frames to the destination endpoint(s) programmed in its settings. If there is no active client listening on that endpoint yet, then send attempts may fail. As soon as a client is activated, it starts to receive frames sent by the server.

UDP_S allows simpler and more restrictive firewall rules because of its unidirectional characteristic. The only traffic that should be allowed is UDP datagrams from the server to a given destination endpoint. All other incoming and outgoing traffic is blocked. Because clients need not know the address of the server in order to initiate a data stream, the server can be made invisible to the external network. Also, because IEEE C37.118 does not have any built-in authentication mechanisms, it is vulnerable to spoofing attacks. An attacker could send a “stop data” command with the source IP address spoofed to be the same as a trusted client, turning off the synchrophasor data stream. Because UDP_S does not use command frames, such an attack is not possible. To keep UDP_S data confidential and allow the integrity of the data to

be checked, transport the datagrams through a VPN to their destination.

Some applications may require sending synchrophasor data bidirectionally between two locations. For such applications, securing the communications via link encryption or VPN would be sufficient, and UDP_S is not necessarily needed. If both can be used, the cybersecurity will be stronger.

2) Remote Access

Accessing the PMU settings remotely through an unsecured protocol such as Telnet poses cybersecurity risks. Because Telnet transports data in a cleartext format, eavesdropping on the network traffic may expose the device settings and login information. This information can then be used for accessing and modifying the PMU settings. Therefore, we recommend that remote access be accomplished through a secured channel, such as Secure Shell (SSH). If the PMU does not directly support such secured channels, then remote access should be accomplished through another intermediary device that supports secured channels. For example, the PMU can be connected to a communications processor using an EIA-232 connection or Ethernet through a firewall that limits the PMU Telnet access to the communications processor. In order to remotely access the PMU settings, engineers can connect to the communications processor over a secured SSH tunnel.

3) Other PMU Ports and Services

PMU devices may have other ports and services in addition to synchrophasors. In order to minimize the attack surface of the PMUs, we recommend that only necessary ports and services be enabled. It is a good practice to disable unused ports and services such as Telnet, File Transfer Protocol (FTP), and Hypertext Transfer Protocol (HTTP). If these services need to be enabled, it is more secure to have external clients access the PMUs through an intermediary device, such as a communications processor, rather than exposing these services through Ethernet.

B. Information Security

A cyberattack can be aimed at the synchrophasor data coming out of the substation, rather than at the substation itself. Because synchrophasor data are used in power system monitoring and control, a potential attack on these data can be just as dangerous. Therefore, securely transporting data is just as important as being able to take accurate measurements from the power system.

We look at three security aspects of synchrophasor data:

- Confidentiality
- Integrity
- Availability

Confidentiality means the data cannot be seen by unauthorized entities. Generally, synchrophasor data are considered confidential for security and competitive reasons.

Integrity means received data are identical to what was sent by the original source and were not modified during transport. Modifying the power system data could cause operators or automated systems to take inappropriate corrective actions or fail to take a correct action. For example, an attacker could modify the voltage angles fed into a wide-area synchronism-check system by adding pi radians to the real value, which could close the breaker while systems are out of synchronism and cause significant damage to the power system equipment.

Availability means synchrophasor data are measured and delivered to the entities that need them in a timely manner. An attacker may interfere with this process in hopes of causing an adverse effect on the system or hiding another type of attack on the system.

These three security aspects of synchrophasor data must be enforced end to end, starting from the PMUs, through the substation network, through the wide-area networks (WANs), all the way to the end-user application. If security is compromised at any time, then the system security is deemed ineffective. In this sense, part of information security is accomplished by substation security practices. Similarly, the security of the control center and the corporate network plays a large role in information security. The security of control centers and corporate networks is outside the scope of this paper, because utilities already have security management in place for them. Here, we focus on information security during the transport between substations and control centers.

As mentioned previously, synchrophasor data are usually transported through IP networks or through a direct serial link, such as EIA-232. Substations are connected to each other and control centers through a variety of communications links. These links may or may not be under the control of the utility. In most cases, we recommend treating these links as untrusted links, although different types of links have different levels of risk. For example, a dedicated fiber-optic link has much less risk of a cyberattack than sending data across the Internet (however, we should never believe that a link has zero risk).

In order to protect data confidentiality and integrity, it is best to encrypt the data during transport across these untrusted links. The encryption may be at the link layer or at the IP layer.

At the link layer, there are various encryption methods available, depending on the type of communications link used. For example, we can use a serial encryption device for serial channels (e.g., EIA-232) [6]. Some wireless radios also provide encrypted channels. The latest SONET multiplexer devices offer hardware encryption at the SONET level, providing levels of security similar to end-to-end encrypted serial links, but with much greater data rates (e.g., Optical Carrier 12 [OC-12], OC-48).

At the IP layer, a VPN can be established to encrypt and secure the synchrophasor data from the rest of the network routable traffic. The termination points of the VPN tunnel are the security gateway that controls the substation access point and a security gateway or another device that supports the same VPN protocol at the control center. In other words, synchrophasor data are encrypted by the security gateway and flow across the untrusted network in an encrypted format. Then the data are decrypted at the VPN termination point at the control center and delivered to the clients. In this way,

PMUs, PDCs, or the clients do not need to support encryption in order to secure the data.

Last but not least, we must keep in mind the potential attacks that try to affect the availability of synchrophasor data. These types of attacks are called denial-of-service (DoS) attacks. Unfortunately, DoS attacks are generally easier to perform than other types of attacks, and at the same time, they are harder to prevent. They usually involve flooding the target device with traffic to consume the target resources and reduce its ability to perform its other key tasks. They can also be accomplished by sending jamming signals to wireless receivers.

DoS attacks can be minimized by making it more difficult for the attackers to reach the targets. Potential targets of a DoS attack are the substation or control center access points (i.e., security gateways) or the devices in the communications links between the access points. For example, if security gateways are linked through a VPN tunnel over the Internet, these gateways are exposed to a very high number of potential attackers. Firewall and VPN mechanisms protect the substation network, as well as the confidentiality and integrity of the synchrophasor data. However, a DoS attack involving a high number of attack points can put the gateway out of service so the synchrophasor data cannot be sent out of the substation. Therefore, using the Internet to transport synchrophasors is not recommended. Wireless links can also be broken using jamming equipment, although the probability and the risk of such an attack should be considered within the context of the specific application. In some applications and locations, we may deem this risk tolerable. Some organizations privately own and manage their own SONETs. Because these networks are physically isolated from external entities, they carry much less risk of an external cyberattack, including DoS and other types of attacks.

Similarly, the GPS signal needed for accurate measurements may be unavailable due to various reasons, such as jammer attacks, wiring issues, or atmospheric conditions. Newer solutions, soon to come to market, will allow the distribution of time signals over a wide area so that an issue with a single GPS clock does not affect the availability of synchrophasor data.

In addition to the possibility of a DoS attack, synchrophasor communications or GPS signals may be unavailable due to other reasons. Putting monitoring and notification mechanisms in place is necessary in order to investigate the cause of system issues and take corrective actions. Also, it is a good practice to account for the unavailability of valid data when using the data in logic and calculations. For example, the logic operations should include checks for IEEE C37.118 time-quality fields in the data frame header and the PMU status fields before taking action based on that data.

Table I summarizes the security concerns and the solutions available for mitigating them.

TABLE I
CYBERSECURITY CONCERNS AND SOLUTIONS

Concerns	Solutions
Unauthorized access to the substation network	Multilayered security and network segmentation
	Firewalls and VPNs
	UDP_S
	Disabled unused ports and services
	Secured engineering access via communications processors
Confidentiality	VPNs or link encryption
Integrity	VPNs or link encryption
Availability	Limited exposure of communications links
	Real-time status and notifications
	Logic checks for data availability and validity
	Wide-area time distribution

IV. EXPERIMENTS AND RESULTS

We performed experiments to illustrate some of the security practices identified in this paper, as well as to verify that these practices do not adversely affect the synchrophasor system performance.

A. Network Latency and Bandwidth

The security practices we discussed in Section III do not affect the measurement or content of synchrophasor data. However, the latency and bandwidth of the communications channels can be affected by VPN tunneling.

We performed latency and bandwidth tests on VPN-capable gateways and routers from four different vendors. Our testing showed that these devices generally provide sufficient bandwidth for synchrophasor communications. Also, the additional latency of VPN tunneling is small enough for most applications. However, these statements may not be true for all applications. The latency and bandwidth requirements of specific synchrophasor applications should be considered and compared with the performance of the network equipment chosen for VPN tunneling.

As we can see in the latency test results in Table II, most gateways seem to have low latencies (< 5 milliseconds) when a small portion of the total bandwidth is in use. However, the results in Table III show that some products have a large increase in latency when we load the device with traffic to use the full bandwidth. Also, we observed an increase in latency as more VPN tunnels were established.

When multiple VPN tunnels used bandwidth, the total available bandwidth was divided across each VPN (see Table IV). All products provide sufficient bandwidth for synchrophasors at a substation where the number of PMUs is relatively low, if the number of VPN tunnels is also low (see Table V). However, some of the products may not be appropriate for deployment at a control center where there is a need to establish VPNs with multiple substations while several synchrophasor streams are coming in.

TABLE II
IDLE LATENCY (IN MS)

Product	Number of VPNs									
	1	2	3	4	5	6	7	8	9	10
Vendor 1	1.913	2.144	1.764	2.647	2.636	3.296	6.232	2.579	4.027	3.538
Vendor 2	1.472	1.82	1.585	1.607	1.717	2.195	2.046	1.691	1.661	1.528
Vendor 3	3.029	2.495	2.807	3.282	3.731	3.014	2.836	3.23	2.497	3.205
Vendor 4	4.672	7.407	5.203	4.412	4.526	5.059	4.734	4.776	4.874	8.918

TABLE III
LOADED LATENCY (IN MS)

Product	Number of VPNs									
	1	2	3	4	5	6	7	8	9	10
Vendor 1	7.765	14.73	19.97	26.46	33.83	32.86	42.71	42.38	48.36	49.21
Vendor 2	9.542	13.28	10.21	12.55	11.91	12.42	14.46	15.92	13.99	13.95
Vendor 3	86.01	93.23	84.9	93.53	95.01	94.19	96.32	96.43	93.9	96.68
Vendor 4	131.3	204.2	195.4	182.7	246.6	203.6	105.5	126.4	120.6	85.4

TABLE IV
BANDWIDTH PER VPN (IN MBPS)

Product	Number of VPNs									
	1	2	3	4	5	6	7	8	9	10
Vendor 1	80.07	38.5	25.5	19.23	15.37	12.79	11.01	9.68	8.84	7.959
Vendor 2	89.62	44.77	29.89	22.43	17.91	14.93	12.84	11.19	9.988	8.994
Vendor 3	6.572	3.273	2.17	1.614	1.281	1.062	0.908	0.791	0.696	0.622
Vendor 4	3.654	1.814	1.203	0.899	0.712	0.597	0.499	0.422	0.372	0.343

TABLE V
TOTAL BANDWIDTH (IN MBPS)

Product	Number of VPNs									
	1	2	3	4	5	6	7	8	9	10
Vendor 1	80.07	77	76.5	76.91	76.87	76.76	77.07	77.44	79.56	79.59
Vendor 2	89.62	89.54	89.66	89.71	89.56	89.58	89.85	89.54	89.89	89.94
Vendor 3	6.572	6.546	6.511	6.456	6.405	6.375	6.355	6.326	6.26	6.215
Vendor 4	3.654	3.628	3.608	3.596	3.561	3.581	3.491	3.372	3.346	3.426

B. Firewall and VPN

1) Test Setup

We set up the test shown in Fig. 5 for demonstrating the additional security that firewalls and VPNs provide.

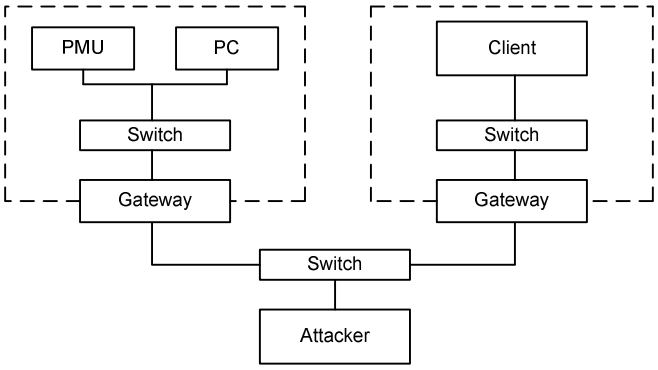

Fig. 5. Test setup

In this setup, we simulate a very simple substation network with a single PMU and a substation PC being protected by a security gateway. On the other side, we have another network that simulates a control center with a client computer receiving synchrophasor data from the PMU.

The two gateways are connected through a switch that simulates an untrusted network. We connected another PC that simulates the attacker on the untrusted network.

2) Configuration

We configured the substation firewall rules to only allow traffic from the PMU to the client at the control center. No other traffic was allowed. We also established a VPN tunnel between the two gateways. We were able to enable or disable the firewall and VPN functions for different test cases.

3) Test Cases

We ran two test cases: a network scan and a man-in-the-middle attack.

We scanned the substation network using a network mapping tool called Zenmap. This tool finds visible IP addresses in a given network and then scans those hosts for open ports and services. It can also graphically display the network topology. We ran this test both with and without the firewall enabled.

When the firewall was enabled, Zenmap was not able to find any hosts. When the firewall was disabled, Zenmap was able to find the two hosts behind the gateway and show their open ports and services (see Fig. 6). It was also able to show the network topology of these hosts (see Fig. 7). An attacker could use this information for focusing attacks on the known vulnerabilities of these ports and services. For example, one of the open ports is TCP 23, which is the default port for Telnet. An attacker might try to sniff traffic going to this port in hopes of intercepting a user login and password.

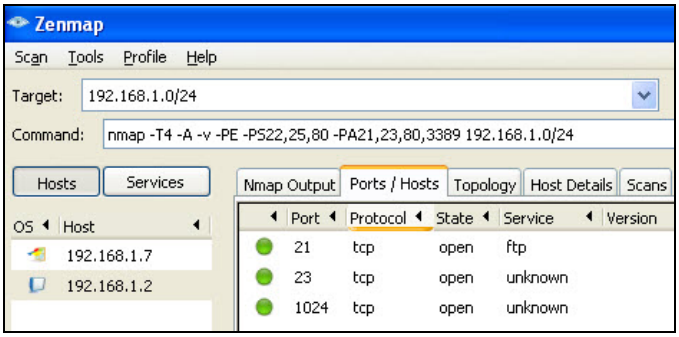

Fig. 6. Zenmap scan results when firewall is disabled (ports and services found)

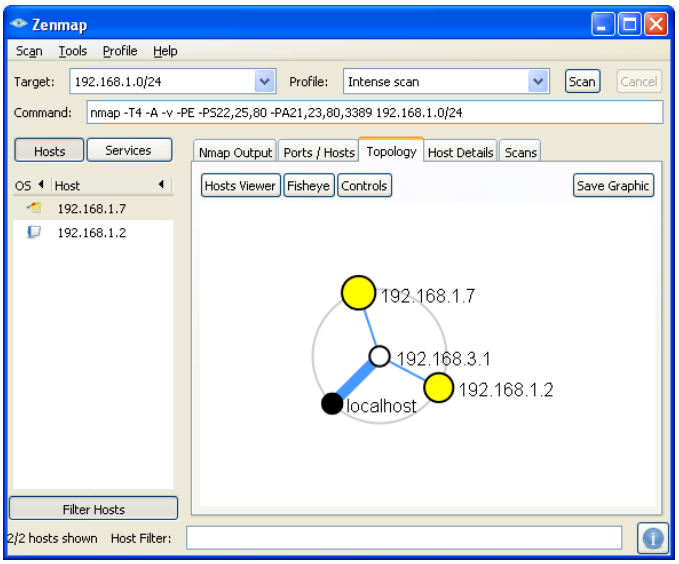

Fig. 7. Zenmap scan results when firewall is disabled (network topology discovered)

Using the attacker computer on the untrusted network from Fig. 5, we performed a man-in-the-middle attack by directing the network traffic to the attacker computer using a technique called Address Resolution Protocol (ARP) cache poisoning. This way, the attacker computer receives all of the traffic between the simulated substation and the control center. The attacker can view or modify the data and then send the data to the recipient as if they were coming from the original sender. For this test, we captured the traffic and tried to view the contents. Also, we replayed the captured traffic and checked to see if the client computer at the control center received it.

We repeated these tests when the VPN was enabled and disabled. When the VPN tunnel was disabled, we were able to redirect and capture the traffic between the two gateways (see

Fig. 8). We were also able to view the contents of the data sent from the PMU because they were not encrypted. Then we were able to replay that data to the client, and the client received the data as if they were sent by the PMU.

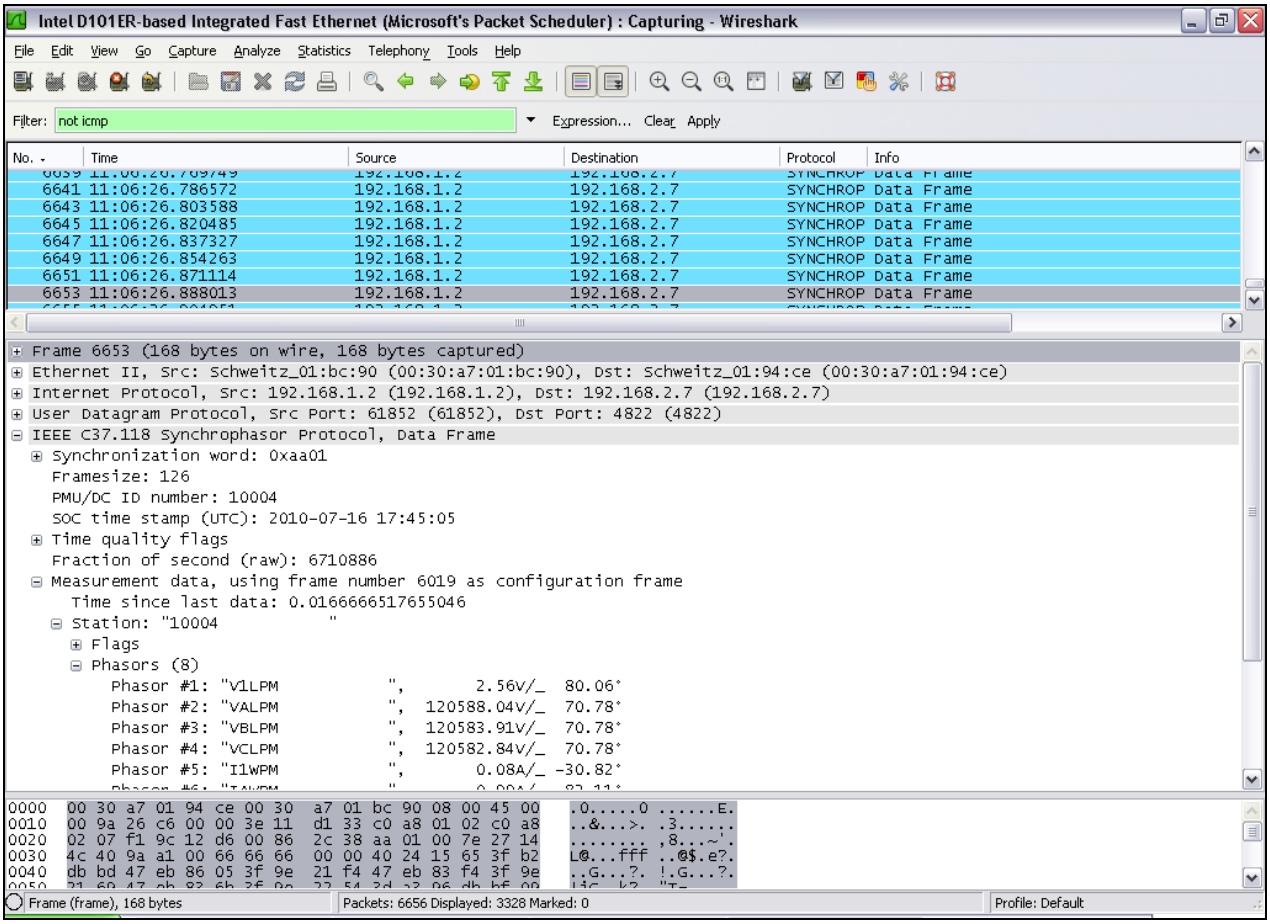

Fig. 8. Synchrophasor data captured during the man-in-the-middle attack (displayed on Wireshark® software)

When the VPN tunnel was enabled, we were still able to capture the traffic between the gateways, but it was encrypted, so we were not able to see the contents or any of the UDP header information that contains the original sender and destination addresses. We were able to pass the encrypted packets through to the destination gateway only once. Replaying the encrypted traffic at a later time did not result in the client receiving the data again, because the gateway checks for a sequence number in the packet after decrypting the packet.

While the VPN tunnel was enabled, we also tried to replay the unencrypted data from the previous test case. The security gateway blocked this traffic as well because it does not allow unencrypted traffic when it knows that those endpoints are in an enabled VPN.

Table VI shows the summary of firewall and VPN test results.

TABLE VI
SUMMARY OF TEST RESULTS

Attack Case	Firewall Enabled	VPN Enabled	Attack Success
Network scan for discovering hosts, ports, and services	No	No	Yes
	Yes	No	No
Man-in-the-middle attack to view and modify data	Yes	No	Yes
	Yes	Yes	No

V. Conclusions

Synchrophasor systems are becoming more critical for managing the power grid, which is growing more complex every day. In the near future, synchrophasor systems are likely to become critical for the health of the grid. Attacks on synchrophasor systems can be as dangerous as attacks on the extant SCADA systems and energy management systems (EMSs). To minimize the risk of security vulnerabilities, synchrophasors should be treated as critical from the ground up.

When implementing synchrophasor systems, the cybersecurity of the existing substation infrastructure should be considered, as well as the confidentiality, integrity, and availability of synchrophasor data.

There are best practices and tools available for securing synchrophasor systems. Having a multilayered architecture protected by security gateways that provide firewall and VPN functions minimizes the risk of an external attacker accessing devices directly connected to the power system. Unidirectional synchrophasor streams and PDCs help define more restrictive firewall rules. Disabling unused ports and tunneling engineering access through an encrypted channel also help protect the substation infrastructure.

Communications links between substations and control centers should be considered as untrusted links in most cases. We recommend encrypting synchrophasor data during their transport through these links.

DoS attacks are harder to prevent. We recommend limiting the exposure of substation and control center access points to potential attackers. Having privately owned networks, such as SONET rings, helps minimize exposure to DoS attacks. If such an attack does occur, it is important to have mechanisms in place for monitoring data availability. We suggest that the client applications and control schemes be designed to handle data unavailability situations.

VI. References

[1] North American SynchroPhasor Initiative. Available: http://www.naspi.org.

[2] A. G. Phadke, "Synchronized Phasor Measurements – A Historical Overview," IEEE/PES Transmission and Distribution Conference and Exhibition 2002: Asia Pacific, October 2002.

[3] E. O. Schweitzer, III, D. Whitehead, A. Guzmán, Y. Gong, and M. Donolo, "Advanced Real-Time Synchrophasor Applications," proceedings of the 35th Annual Western Protective Relay Conference, Spokane, WA, October 2008.

[4] A. Johnson, R. Tucker, T. Tran, J. Paserba, D. Sullivan, C. Anderson, and D. Whitehead, "Static Var Compensation Controlled Via Synchrophasors," proceedings of the 34th Annual Western Protective Relay Conference, Spokane, WA, October 2007.

[5] IEEE Standard for Synchrophasors and Power Systems, IEEE C37.118-2005, 2006.

[6] H. J. Altuve Ferrer and E. O. Schweitzer, III (eds.), *Modern Solutions for Protection, Control, and Monitoring of Electric Power Systems*. Schweitzer Engineering Laboratories, Inc., Pullman, WA, 2010.

VII. Biographies

John Stewart is a principal engineer in the architecture group of the Tennessee Valley Authority Transmission Power Control Systems organization. In his more than ten years of power control systems and communications experience, he has been the technical lead for multiple infrastructure and pilot projects focused on key areas of technology associated with grid data, transport, and control. He is currently leading the access control technology team in response to the NERC CIP standards within the transmission organization. John holds a BS in electrical engineering with a focus on telecommunications systems from Tennessee Technological University.

Thomas Maufer, published author of three books for Prentice-Hall, is the Director of Technical Marketing for Mu Dynamics. He has managed local-, metropolitan-, and wide-area networks at NASA/GSFC, participated in the InteropNet Network Operations Center Team, and worked in Silicon Valley, building all types of Ethernet- and IP-based networking products, from semiconductors to switches and routers, now including test equipment. Much of his career has centered around security, both in products and in protocols. He has been fortunate to have spent many years contributing to IEEE and IETF standards, including IEEE 802.11e/g/i, IEEE 802.1X, IPsec, IPv6, OSPF, and others. As of this writing, he has contributed to 29 United States patents. When he is not helping to design protocols or the products that implement or test them, he is relaxing with a single-malt Scotch.

Rhett Smith is a development manager for the security solutions group at Schweitzer Engineering Laboratories, Inc. In 2000, he received his BS degree in electronics engineering technology, graduating with honors. Rhett is working on two U.S. Department of Energy control system security cooperative agreements. He is the project director for the Hallmark project and is one of the principal investigators on the Lemnos project. Rhett has his GSEC, GIAC Security Essentials Certification, and is a Certified Information Systems Security Professional (CISSP).

Chris Anderson has an AAS from ITT Technical Institute and is currently working on his BS in electronics engineering technology through DeVry University. He has been with Schweitzer Engineering Laboratories, Inc. (SEL) since July 1999. He worked in product development for transmission protection products and then as an associate product engineer supporting transmission protection before becoming responsible for the development and support of SEL synchrophasor measurement technology. In February 2008, Chris transferred to the field as an integration application engineer for the southeastern United States, where he currently makes his home in Dothan, Alabama. He has been involved with the Eastern Interconnection Phasor Project (now the North American SynchroPhasor Initiative) since 2004 and with the introduction and implementation of synchrophasor technology at utilities spanning the United States and overseas.

Eren Ersonmez is an integration and automation engineer at Schweitzer Engineering Laboratories, Inc. He focuses on communications and security of distributed synchrophasor systems. Eren participates in the NASPInet development efforts within the NASPI Data and Network Management Task Team. He previously worked as a software engineer for a major semiconductor manufacturer. He received his BS in information systems from the University of Idaho.

© 2010 by Tennessee Valley Authority, Mu Dynamics, Inc., and Schweitzer Engineering Laboratories, Inc.
All rights reserved.
20100913 • TP6449-01