How SDN Can Improve Cybersecurity in OT Networks

Colin Gray
Schweitzer Engineering Laboratories, Inc.

Presented at
IEEE ROC&C 2018/2019
Acapulco, Mexico
March 6–8, 2019

Originally presented at the
22nd Conference of the Electric Power Supply Industry, September 2018

HOW SDN CAN IMPROVE CYBERSECURITY IN OT NETWORKS

Colin Gray
Schweitzer Engineering Laboratories, Inc.
3 Hampstead Way
Hamilton, New Zealand
Email: papers@selinc.com

***Abstract*—With the business need for interconnectivity between information technology (IT) and operational technology (OT) networks becoming more prevalent, cybersecurity is now a critical aspect of how that interconnectivity is achieved. This paper discusses the differences between implementing cybersecurity measures in IT and OT networks, introduces software-defined networking (SDN) technology, and describes how SDN can be used to build a robust and secure OT network. This paper is intended for engineers interested in secure network design strategies and provides a platform for discussion about what SDN is and how it can be implemented in an OT network to meet cybersecurity, performance, and management needs.**

I. INTRODUCTION

In the last decade, significant advancements have been made in how industrial control systems (ICSs) monitor, control, and manage an organization's infrastructure. Implementation of substation-hardened Ethernet local-area networks (LANs) has paved the way for more precise and sophisticated tools for retrieving and analyzing critical real-time system information. The advanced capabilities of today's modern digital measurement and control intelligent electronic devices (IEDs), coupled with the implementation of Ethernet access directly from the device, has opened the door for Transmission Control Protocol/Internet Protocol (TCP/IP) services to access an unprecedented amount of data from these devices.

IEC 61850 provides high-speed peer-to-peer communications between IEDs, enabling faster and more accurate control and restoration in the event of network disturbances. IEEE C37.118 synchrophasors and IEC 61850-9-2 Sampled Values provide high-speed sampling of real-time data directly from the secondary source with millisecond accuracy.

While this information has traditionally been firmly in the domain of the operational technology (OT) network, a wider audience has seen the need to bridge the gap between OT and information technology (IT) networks. While IEC 61850 and other Ethernet-based supervisory control and data acquisition (SCADA) protocols operate at the machine-to-machine (M2M) level, a large portion of network activity now involves person-to-machine (P2M) interconnectivity [1].

The demand for access to these data is skyrocketing as more organizational departments use an array of services for a variety of purposes, such as for analyzing post-event disturbance records, accessing metering data, and gaining secure access directly to a device for diagnostics and troubleshooting to ensure that the highest possible service availability is maintained. This, of course, inevitably raises the specter of cybersecurity and how the traditional Ethernet network deals with securing critical infrastructure information. The fifth version of the North American Electric Reliability Corporation Critical Infrastructure Protection (NERC CIP) reliability standards (known as CIP Version 5) outlines a set of requirements for securing the traditional LAN from unwanted intrusions and attacks. The impact of such attacks has been well documented in the past.

Recent research has shown that software-defined networking (SDN) can provide far better packet delivery performance, greater cybersecurity, and greater centralized seamless control than can a traditional Ethernet LAN. SDN has revolutionized the way IT system managers program their networks to keep up with a vastly diverse range of users and applications. This paper introduces SDN and how it can be implemented into an OT network to provide a greater level of security and performance.

II. WHAT IS SDN?

Ethernet technology is segregated into two distinct service planes: the control plane and the data plane. The control plane is responsible for determining which network path a packet should use to reach its destination. The data plane is responsible for carrying out the control plane requirements by determining the packet contents and informing the individual switches of how to forward the packet around the network [2].

The key difference that SDN introduces is the separation of the control plane and data plane. In a traditional network as shown in Fig. 1(a), a managed network switch uses Spanning Tree Algorithms (STAs) via the Rapid Spanning Tree Protocol (RSTP) to determine a single path connection between two Layer 2 switches or middleboxes. This switch is also responsible for reconverging the network after a fault is detected. A predetermined root bridge periodically sends special data frames called Bridge Protocol Data Units (BPDUs) to determine the state of the network paths and the appropriate action to take in the event of a change in the network. Embedded in the Ethernet frame to manage throughput reliability are other network control parameters, such as IEEE 802.1Q virtual LAN (VLAN) tags and IEEE 802.1p Class of Service tags. Each middlebox must be able to internally manage the control plane logic and then execute the necessary data-forwarding requirements in the data plane logic.

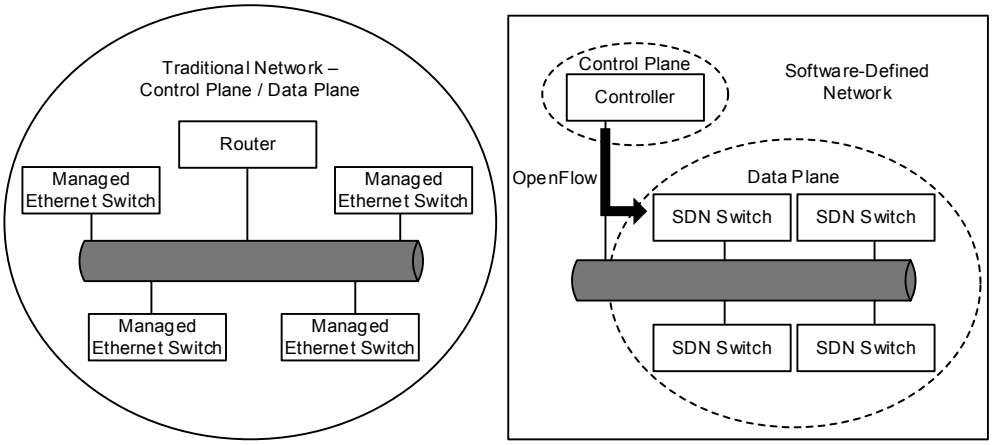

Fig. 1. Traditional LAN Architecture (a) and SDN Architecture (b)

SDN, on the other hand, implements a logical separation of the control plane and data plane, as shown in Fig. 1(b). It does this by removing the control plane logic from the switch and introducing a network controller, which determines the data flow criteria and defines a set of control plane flow rules that are sent to the SDN switch. The SDN switch then simply executes the flow rule as predetermined by the controller [3]. SDN gives system administrators much greater control over how specific data are forwarded around the network and allows network configuration to be designed to suit the data exchange needs of the application. An SDN network, therefore, does not need to rely on STAs and RSTP to determine a best-path connection between two devices or wait for an RSTP reconvergence decision after a failure in the network before re-establishing a new data connection. This means that SDN can program more than one active connection to an IED at the same time and create redundant data paths in the same network.

SDN uses the OpenFlow protocol as the interface between the network action required for data packets (in the control plane) and the hardware that carries out that action (in the data plane). OpenFlow passes flow rules from an OpenFlow controller to OpenFlow-enabled SDN hardware. The traditional network switch must make its own decision on how data flows based on network conditions. Because flow rules have been predetermined in the data plane OpenFlow-enabled hardware, there is no need for the data plane hardware to make these decisions. The network simply follows its flow table match rules to decide what to do next.

III. IT SDN vs. OT SDN

The U.S. Department of Energy recently sponsored two research projects to determine what advantages SDN could offer OT networks. It is important to note that while the SDN technology applied to OT networks has been altered to fit industry-specific standards and requirements, the underlying SDN standards applied to the IT SDN network did not have to change, paving the way for rapid development. The research project results far exceeded the expected performance and security criteria set by the research participants and helped set the benchmark for OT SDN deployment. Table I summarizes the differences between how SDN is applied to an OT network versus an IT network.

TABLE I
OT SDN VS. IT SDN [3]

Key Attribute	OT SDN	IT SDN
Network state	Persistent	Dynamic
Network control	Purpose-engineered	Traffic reactive
Controller purpose after switch deployment	Monitor	Control
Security	Deny-by-default	Plug-and-play
Fault healing speed	Link detect	Flow setup time
Network management	Traffic reactive	Fault reactive

Purpose-engineered network control has significant advantages over IT traffic reactive control. Generally, an IT middlebox will be plug-and-play, which is great for ease of installation because the system manager does not have to configure many parameters. However, it makes network changes more complex because the ultimate decision about which path a packet will take is determined by the STA and RTSP. Plug-and-play networks always attempt to deliver a packet to its intended destination. In many cases, such as when using Address Resolution Protocol (ARP), the networks also attempt to resolve the packet query by broadcasting to all hosts in the subnet. This practice has obvious security ramifications.

Purpose engineering gives the system engineer complete control over the desired path a packet should take during normal and abnormal network conditions. While some argue that purpose engineering is inherently more difficult to configure because of the deny-by-default nature of the network, it is the only way we can be sure the critical data in networks can be delivered to the intended destinations no matter what state the network is in.

IV. CYBERSECURITY BEST PRACTICES

The following list provides best practices for effective cybersecurity (this section previously appeared in [4]):

- Know the system (NERC CIP-002).
- Baseline approved systems (NERC CIP-010).
- Practice need-to-know policies (NERC CIP-004).
- Train staff (NERC CIP-004).
- Establish a defense-in-depth architecture (NERC CIP-005 and 007).
- Protect the data (NERC CIP-011).
- Log and monitor the system (throughout NERC CIP).
- Have redundant communication paths (throughout NERC-CIP).
- Maintain peak performance of the system (throughout NERC CIP).
- Establish access controls (NERC CIP-005).
- Plan and practice for incidents and responses (NERC CIP-008 and CIP-009).
- Practice physical security (NERC CIP-006 and CIP-014).

Whether designing an IT or OT network using traditional Ethernet or SDN, network designers need to know and understand what they are trying to achieve in any cybersecurity effort. The bottom line is that we want to preserve the integrity of the system we are protecting so that it can operate in a safe, reliable, and economical manner.

First, we must get to know the system we are tasked with protecting. Second, we need to baseline the approved systems. We should also practice need-to-know policies and only share information about the security systems with those whose jobs require that information. The principle of least privilege allows authorized people to gain just enough access to the system to accomplish their jobs. Cybersecurity success hinges on people performing their jobs in a secure way, and the only way people will know how to do this is if they are trained.

Next, we need to engineer a defense-in-depth architecture with multiple layers of complementary defensive technology that block or slow down attackers so that there is time for an operator to be alerted and respond to the attack. Data protection should apply to both data in transport and at rest. This usually means that cryptographic solutions should be used, including encryption, digital signatures, and authentication. Another solution to data protection is offline storage with removed and turned-off memory.

We need to continually monitor what is happening on the system so we can respond as quickly as possible. Redundant communications paths are key to doing this. When one path is under attack or rendered unavailable, redundant paths allow the logs and alerts to get out and an incident response team to get in. We must also maintain peak performance of the system; maintenance of the devices and technology that make up the system includes hardware tests, software patch management, and configuration validations. Additionally, we must ensure that we have appropriate access control in place, limiting logical access to only those people who are authorized to gain access and making sure we have proven who they are before access is granted. This is authentication.

Cybersecurity incidents will happen, so we must plan for them, practice appropriate responses through proven policies and procedures, and provide physical security measures.

V. Traditional IT vs. OT Network Cybersecurity

IT cybersecurity has clearly established security procedures, tools, and applications. OT products are not typically tested with antivirus solutions because their virus databases need constant updates. It is rare for an OT network to include security logging, incident response plans, or forensic analysis. Once an OT computer is installed and put into service, there are typically very few updates or changes. Enterprise computers are generally replaced every 2 to 3 years, but it is normal for an OT computer to be in operation for 10 to 20 years. Because of this, it is likely that new vulnerabilities on an OT network will remain unpatched until the control system is replaced.

Blacklisting and whitelisting are common techniques used to identify communications or actions on a network and either restrict them or allow them to occur. A blacklist contains everything that should be denied, and allows anything that is not specifically denied. Blacklisting is commonly used in enterprise networks to block access to known "bad" elements, such as infected websites or compromised IP addresses. In contrast, a whitelist contains a complete set of allowable actions. Whitelists, which block anything except what has been approved, are common in ICS applications. An ICS has a much smaller set of possible communications or actions, which is unlikely to change. It is possible to create a list of exactly what communications or actions should be taking place and to deny everything else in an ICS, but this would be nearly impossible in an enterprise network.

Table II highlights some of the key differences in implementing cybersecurity measures in an IT network versus an OT network.

Table II
Cybersecurity Measures Implemented in IT vs. OT Networks [4]

Application	IT Network (Enterprise)	OT Network (ICS)
Antivirus	Very common; easily deployed and updated; blacklisting used	Restrictions on new systems; difficult to deploy on legacy systems; whitelisting used
Patch management	Easily defined and automated; remotely deployed enterprise-wide	Typically requires vendor validation and owner/operator testing
People	Office environment	Operations environment
Incident response and analysis	Well-defined, understood, and deployed; extensive analytical tools available	Not common after system restoration; no in-depth analysis beyond event recreation
Asset management	Performed periodically	Performed infrequently
Cybersecurity testing and auditing methods	Can use widely available tools and methods	IT tools and methods not suitable for ICS networks
Technology lifecycle	2 to 3 years	10 to 20 years
Software changes	Frequent	Rare

VI. OT SDN CYBERSECURITY

Cybersecurity is critical to the reliable operation of ICS networks. The same aspects found in the cybersecurity best practices apply to SDN with organizational policies, both with the procedures to be followed by people interacting with ICS network devices and with the implementation of the devices themselves. SDN technology leverages the core goal criteria to provide business benefits, situational awareness, and incident response planning while maintaining an architecture that is simple to implement. Applying the components of the cybersecurity best practices to SDN produces a robust and reliable ICS network capable of maintaining the highest levels of security and able to respond to varying threat levels through proactive network engineering.

A. Know the System

A solid understanding of network architecture is developed from the onset when SDN network design is considered up front. Any network changes can be undertaken by following strict change control management policies and allowing the deployment of powerful whitelisting security techniques.

An ICS network typically has known embedded devices that perform specific tasks for an extended period of time. This simplifies the application of security controls, ensures that unknown devices or applications are dropped, and ensures that appropriate measures are taken. This means that the SDN deny-by-default strategy of enabling authorized traffic rather than trying to find unauthorized traffic eliminates the need to constantly manage patches, system updates, and device security rules and policies for deep packet inspection (DPI) devices and intrusion detection systems (IDSs).

B. Baseline Approved Systems

Network designers are continuously striving to find streamlined methods to confirm that the network is operating at the approved baseline. Documenting this, as required by NERC CIP, and ensuring that the network adheres to baseline performance criteria can be difficult in a traditional network. Because these networks rely on many protocols (such as Link Layer Discovery Protocol [LLDP] and Syslog) to collect information and determine the near real-time network state through a variety of dynamically changing broadcast and multicast mechanisms, it is difficult to verify with any certainty that the actual data flow matches the expected baseline.

Guaranteeing the contents of a packet beyond Layer 2 requires additional hardware for DPI. SDN control plane separation has this facility already available. OpenFlow provides an array of counters that can be configured for reporting packet and byte counts for all types of flow entries and can be configured to monitor for bursts of activity. This gives SDN a distinct advantage in its ability to record and maintain a baseline through continuous monitoring in real time.

C. Practice Need-to-Know Policies

SDN deny-by-default access makes it simple to implement need-to-know policies. User access to device-specific activities can be closely monitored. When coupled with SDN management platforms, authentication, authorization, and accounting (AAA) proxy criteria can be strictly enforced.

D. Train Staff

With SDN flow table capabilities, switches can be configured for training scenario modes through path duplication of ICS data without compromising operational integrity. This gives greater access to baseline data and eases the execution of security control measures.

E. Establish a Defense-in-Depth Architecture

For an STA network, a defense-in-depth architecture recommends the implementation of six levels, each with its own security and monitoring requirements. As detailed in Fig. 2, each level is assessed on its need for confidentiality, integrity, and availability (CIA). Level 1 and 2 devices should be separated from the rest of the system as much as possible to reduce their attack surfaces. SDN traffic engineering ensures a greater level of security through direct device-to-device path engineering.

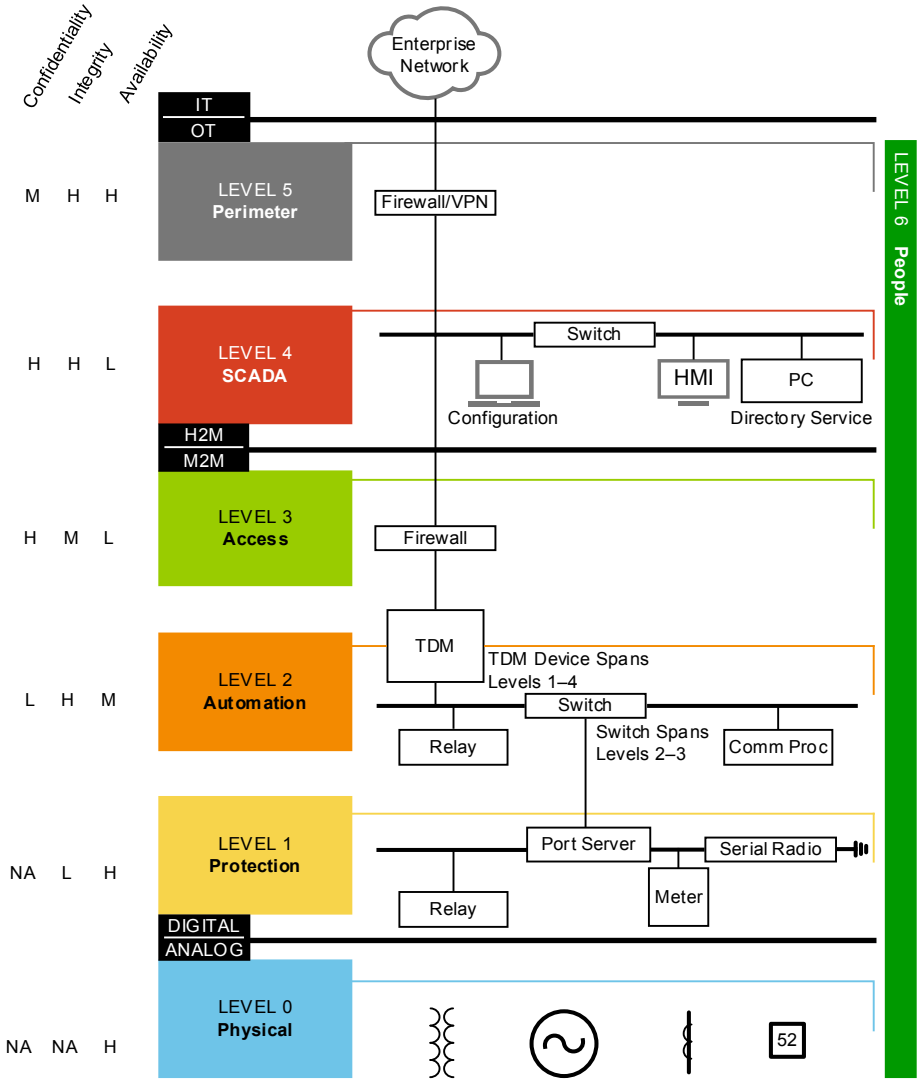

Fig. 2. Defense-in-Depth Architecture

Cyber attackers are always looking to exploit vulnerabilities in an STA network. Media access control (MAC) table flooding attacks can be used to force a switch to refresh its MAC table, which causes a significant portion of the incoming frames to be flooded out all ports, thus creating an opportunity to capture data that reveals legitimate MAC addresses. ARP spoofing (in which attackers attempt to redirect traffic to their IP address by sending ARP messages to the network) or BPDU spoofing (in which unauthorized BPDU frames are sent into a network to alter how packets are forwarded through the network) are both common techniques of intercepting packets for the launching of denial-of-service (DOS) or man-in-the-middle attacks [3].

In the case of the MAC table flooding attack, mitigation is achieved by MAC-locked port technology. ARP spoofing can be mitigated with static ARP tables so that hosts do not need to reply to ARP requests. However, in both instances, maintaining a network that periodically requires new or replacement hardware requires additional resources and funds to unlock and relock MAC addresses and manually update ARP tables. BPDU spoofing requires port monitoring and disabling and BPDU filtering so that any BPDU packets arriving at the port are dropped.

By default, STA switches forward all packets when more information is required so that security relies on the quality of packet filtering. This is required for every individual switch in a network, resulting in more administrative overhead to continually monitor, patch, and maintain the highest level of security.

The deny-by-default architecture of SDN technology ensures any unauthorized packets are dropped and only packets that match proactive traffic-engineered flows are forwarded. MAC table flooding and BPDU spoofing will not work in SDN because SDN switches do not have MAC tables or use BPDU packets. ARP spoofing is overcome by engineering ARP requests and responses between devices.

The same concept of security levels can be implemented in SDN in a much simpler fashion. Fig. 3 shows the implementation of security zones based on the collection of devices required to interact within that zone. This keeps each zone or region isolated with trust established between all the devices within that zone. SDN flow rules to apply simple techniques (such as MAC and IP address masking for traffic external to the security zone and blocking of all multicast and broadcast traffic outside the security zone) provide a powerful tool to manage the security needs of an ICS network.

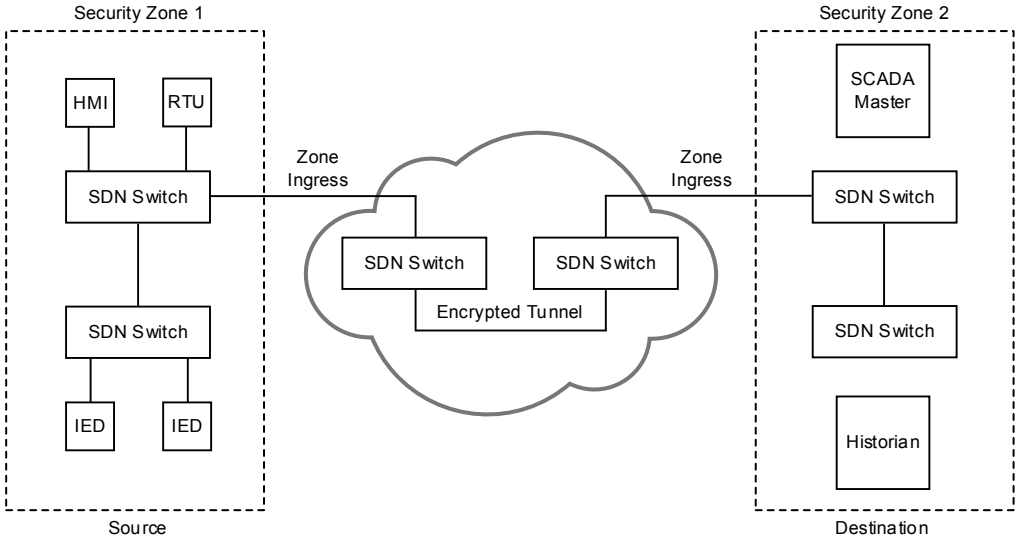

Fig. 3. SDN Security Zones [2]

With available OT SDN cybersecurity software, the security state of the network can be managed through policy implementations that can be changed as the security state of the network changes. This provides security orchestration for the flow controller through complete network visibility, programmable security zones, situational awareness, and security policy management through whitelisting all devices on the network. Fig. 4 details the architecture of third-party applications that are being developed to provide SDN-based cybersecurity policies and measures at the enterprise level.

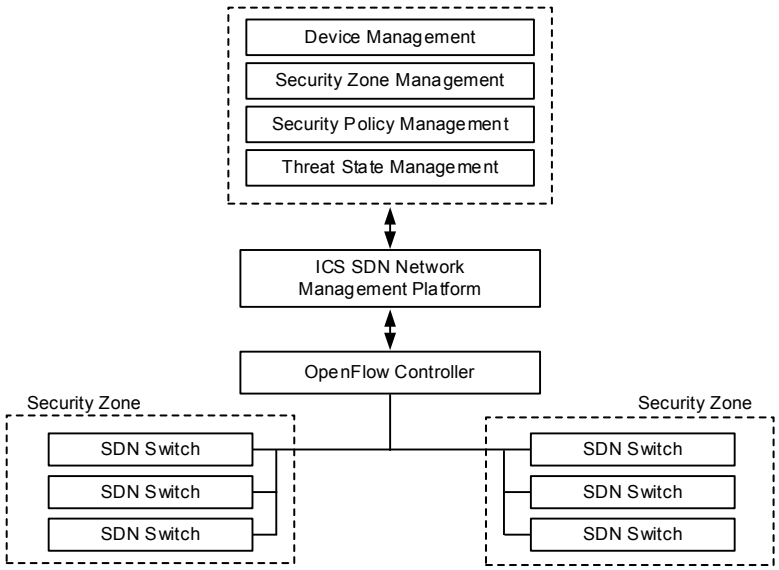

Fig. 4. Application Programming Interface (API) to SDN Cybersecurity [2]

The security management policy allows a network to be built with different operational states and describes how the network is treated in response to changing threat levels. This is a dramatic shift from reactive network analysis, detection, and action to one of proactive, preplanned security controls.

The threat state model can be divided into six levels of trust as determined by the defense-in-depth model and can be configured to decide whether human intervention is required for a request to transition levels or whether it can be done automatically. Each level can be assigned a predetermined security policy and security zone, resulting in a threat-based approach to security controls for the system.

F. Protect the Data

Data protection is paramount in critical infrastructure operation. Archiving OT SDN data is conducted with the same devices used by an OT STA network with the added benefit that SDN flow tables provide greater flexibility in determining which data are archived. This leads to more efficient bandwidth use across the network.

G. Log and Monitor the System

Control plane communications in an STA network are not protected through encryption or authentication. However, OpenFlow communications between the control plane and the data plane are mutually authenticated using X.509 certificate verification and are encrypted with Transport Layer Security (TLS), making configuration communications on the centralized control plane a very small attack surface.

DPI, IDSs, and Intrusion Prevention Systems (IPSs) in an STA network are typically difficult to deploy in an ICS because of the frequent need to update certificates, signatures, and keys, and because of the associated hardware necessary to enforce them. IDS hardware is generally required to inspect all packets and therefore requires a dedicated port for packet forwarding and inspection. This can be a significant amount of traffic that requires hardware not rated for industrial or substation use. Whitelisting is the preferred method for antivirus mitigation in the OT STA network, but this tends to be firmware-embedded and not easily updated.

The proactive traffic-engineered solution of SDN provides the ability to inspect every packet at every hop and at multiple layers. MAC, VLAN, and ARP addresses; traffic types; source IP addresses; destination IP addresses, and associated port numbers can all be inspected and whitelisted as they traverse the network. With this level of multilayer packet inspection though SDN flow matching, only packets that do not meet the flow and match criteria need to be sent to an IDS. Any packet that does not meet an authorized flow entry is sent for DPI. The ability of SDN to inspect packets in this manner greatly reduces the traffic inspection required by a traditional IDS, thus reducing instances of false positives and eliminating the need for specialized hardware, which allows the deployment of industrial-rated hardware.

Because only packets that have not been engineered to be on the network are sent to the IDS, traffic-engineered networks protect against insider threats and mistakes. Either a detected packet should be on the network and a flow should be added to authorize that packet, or the packet is unauthorized and we can determine its origin and how to deal with it.

H. Have Redundant Communications Paths

Redundant communications paths are essential for the reliable operation of an OT network and to provide alternative paths in the event of a failure or attack. An STA network allows only one active path between devices, and in the event of one of these scenarios, it relies on RTSP to reroute packets according to its algorithms via the next least costly path. The network healing time with this technology can be on the order of 50 to 100 milliseconds.

SDN allows the use of more than one active connection on an IED. A typical IEC 61850 scheme on an STA network, for instance, would require the IED to detect a link failure before failing over. SDN allows the engineering of simultaneous active paths, providing the IED the ability to send Generic Object-Oriented Substation Event (GOOSE) messages via two separate paths within the same network. A proactive pre-engineered network can then react to network failures much faster because it only requires the switch to read its flow tables every processing cycle to determine the next course of action to take. Therefore, network healing time is significantly reduced to approximately 100 microseconds or less.

OpenFlow also provides priority-based port grouping as another method for path redundancy. An SDN switch uses the highest priority port to send packets, and if this port is unavailable, the switch uses the next available port. This eliminates the need to resend packets. Even if the port is not part of a group, a packet can be sent back out the ingress port and redirected out of an alternative port. Such options are impossible to achieve in an STA network.

Of course, the most reliable method for redundancy is to provide two separate networks physically isolated from each other in a dual primary redundant network. Parallel Redundancy Protocol (PRP) and High-Availability Seamless Redundancy (HSR) are two such duplication methods. SDN traffic engineering can emulate the benefits of these technologies, although the number of flow rules required may be a limiting factor.

I. *Maintain Peak Performance of the System*

SDN whitelisting ensures that the network performs at peak levels of efficiency and reliability, as no unplanned or unauthorized devices or applications can consume network resources. Baseline statistics and known traffic volumes that do not change dynamically without preprogramed engineering make SDN peak performance adherence and monitoring easier to quantify than they are on an STA network.

J. *Establish Access Controls*

Access controls form part of security policies and security zone plan management. Because SDN is an interoperable Ethernet technology, network hosts that provide user authentication services, such as Lightweight Directory Access Protocol (LDAP), can still provide AAA proxy services to users crossing from the IT network to the OT network.

K. *Plan and Practice for Incident Responses*

ICS networks have well-defined incident response plans to ensure preservation of life and equipment as well as cybersecurity threats. SDN complements this approach through its proactively engineered network plan that allows quick changes to the operational state of the network in microseconds, irrespective of the network size or topology.

L. *Practice Physical Security*

As with an STA network, SDN can be integrated with physical access control systems.

VII. Implementing Packet Sniffing in an SDN Network

DPI, IDSs, and port mirroring, or "packet sniffing" as it is more commonly known, are packet-analysis techniques used to inspect packet contents or sometimes entire data streams of information flowing across the network so that unwanted or malicious content can be detected. In the example cyber attack timeline in Fig. 5, an attacker may have malware sitting undetected on a network for many months. SDN offers significant improvements in packet-sniffing capabilities, allowing detection of intruder activity at the reconnaissance phase.

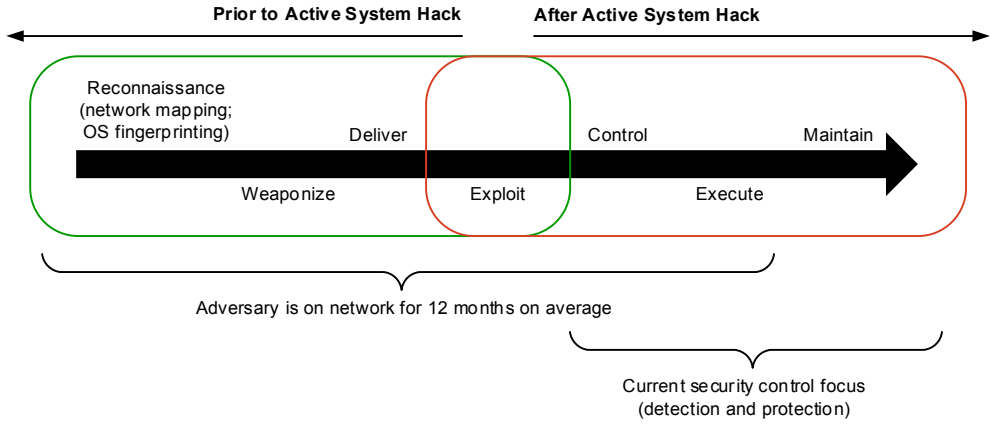

Fig. 5. Cyber Attack Timeline [2]

With SDN flow matches and conditions, we can be more selective with the packets needing further inspection in a switch and redirect them to the appropriate analysis device. A DPI application requires all packets to be rerouted to the DPI analysis device and then possibly return to the originating switch. The application may only be interested in HTTP traffic, for instance, and lets all other traffic traverse the network without additional scrutiny. SDN allows packet identifier tags to be added so that the packet can be traced back to the switch where it originally entered the network, which is valuable to know when the packet is rerouted to the DPI device and is required if the DPI return route is needed.

SDN switches that allow the network engineer to add or remove VLAN IDs determine the switch by which a packet first entered the network. When packets already have a VLAN ID, bit masking can be used in the VLAN field to identify the packets that have been previously marked for inspection and need to be routed to the analysis device. This combination of bit masking and VLAN ID assigning ensures that any packets that were previously sniffed and tagged in a downstream switch are forwarded, will remain isolated from other traffic, and will be ensured a unique VLAN ID as they enter each switch [5]. This also ensures that if the packet is not dropped by the DPI analysis device, it can be forwarded to the original destination if required.

An SDN system assumes that all packets have a pre-engineered match in a flow table and therefore any packet that causes a table miss is either not required or may be an intruder attempting to access the network. For an IDS and port-mirroring analysis, the packet is simply duplicated and sent to its original destination and to a predefined switch port for forwarding to the DPI analysis device.

When SDN packet sniffing is used in conjunction with analytical tools that can respond to the detection of an attacker (who is attempting to map the network through ARP scanning or to identify the operating system (OS) of a network through OS fingerprinting techniques), it becomes very difficult for the attacker to move to the weaponize and deliver phase. These responses could be in the form of false IP addresses in ARP replies or OS information preventing the attacker from mapping a network or identifying the OS.

VIII. Conclusion

Through control plane abstraction, SDN technology removes the restrictions of STA, both in physical network design and operational performance. ICS purpose-engineered networks are redefining the criteria for the creation of new best-known methods of delivering information between services and applications. They are doing this while still improving system performance in technical, policy, and procedural ways without breaking the existing interoperability with Ethernet.

SDN allows network engineers to support even the most demanding applications used to control, operate, and monitor critical infrastructure in an ICS network. Centralized monitoring and change control services reduce the risk of network disruption. The traditional extended wait times for an STA network to converge have become a thing of the past since SDN switches are preprogramed to forward packets in such events.

Although not a completely lossless technology, as packets could still be lost in transit if they are on a link or in a switch that fails or if a port buffer is overrun, OT SDN is effectively a next ingress packet-healing technology. SDN clearly provides significant advantages in cybersecurity over STA, with its core architecture complementing the ICS core attributes and predictable whitelisting technologies. NERC CIP compliance can be achieved with near real-time centralized reporting and monitoring of port activity, traffic volumes, and application status without the need to manually deploy engineers to gather and analyze these data.

SDN challenges the best-known practices for STA network architectures, router design, combined control plane and data plane switch design, and the services they use. The proactive engineering design of SDN provides DPI and an IDS at every switch while simultaneously eliminating the services that increase the attack surface in the STA network. MAC tables and BPDU packets are not used in SDN, and ARP spoofing through broadcast and multicast blocking can be eliminated. In fact, the whole concept of what a subnet is and how IP addresses are grouped can be challenged with flow rules that manage packet flow based on any layer of packet contents.

The exposure to attack and the recovery from a detected threat or intrusion is vastly reduced with the evolution of third-party OT-focused SDN management platforms. Coupled with SDN multilayer capabilities, management platforms that provide corporate-level threat models, security zones, and policies allow network operators to have preplanned contingencies for mitigation of a threat detection. As the future possibilities for OT SDN development and design continue to evolve and a greater understanding of the performance and security aspects of SDN are further understood by asset owners, SDN will become the go-to architecture for making exciting changes in ICS networks.

IX. References

[1] D. Dolezilek, C. Gordon, and D. Anderson, "Applying New Technology to Improve the Performance of Teleprotection Communications," proceedings of the 13th International Conference on Developments in Power System Protection, Edinburgh, United Kingdom, March 2016.

[2] R. Hill and R. Smith, "Purpose-Engineered, Active-Defense Cybersecurity for Industrial Control Systems," August 2017. Available: https://www.selinc.com.

[3] M. Hadley, D. Nicol, and R. Smith, "Software-Defined Networking Redefines Performance for Ethernet Control Systems," proceedings of the Power and Energy Automation Conference, Spokane, WA, March 2017.

[4] C. Gray, "An Exercise in Trust: Examining the Future of Cryptographic Technology in Electrical Control Systems," proceedings of the South East Asia Protection, Automation and Control Conference, Melbourne, Australia, March 2017.

[5] J. Dearien, "Implementing Packet Sniffing (IDS, DPI, Port Mirroring, etc.) in an SDN Network," SEL Application Guide (AG2017-18), 2017. Available: https://www.selinc.com.

20180720 • TP6891-01