A Case Study Assessing the Effects of Cyber Attacks on a River Zonal Dispatcher

Ronald Joseph Wright[1], Ken Keefe[2], Brett Feddersen[2], and William H. Sanders[1]

[1] Department of Electrical and Computer Engineering,
University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA
{wright53, whs}@illinois.edu
[2] Information Trust Institute,
University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA
{kjkeefe, bfeddrsn}@illinois.edu

Abstract. A river zonal dispatcher is a system that sends collected environmental data to a national dispatcher and sends warnings in case of danger (such as flooding of river basins). If the system fails to function normally, warnings may cease, putting lives and property in serious peril. We have examined the security of a river zonal dispatcher using the ADVISE modeling formalism in the Möbius modeling tool. This work both illustrates the usefulness of ADVISE in choosing among alternative approaches to system security and provides a quantitative evaluation of the dispatcher itself. In doing so, it shows whether intrusion detection systems (IDSes) make a difference in the behavior of an adversary, and which path of attack is most attractive to particular types of adversaries.

Keywords: control systems security, quantitative security metrics, state-based security model, discrete event simulation

1 Introduction

Critical infrastructures must be resilient to real-world threats. According to the Department of Homeland Security, dam systems are dependent and interdependent on a multitude of sectors, such as the water sector (for delivering potable water to customers) and the emergency services sector (for delivering water in case of emergencies such as firefighting). Dam systems contribute to numerous projects, such as hydroelectric power generation, navigation control, levees, and waste impoundments [10, 11]. There is a need to be vigilant against any threats to the integrity of such water control systems so that they function with little or no interruption.

Incidents have occurred that have threatened the integrity of water control systems. In 2001, a former employee of a sewage treatment plant in Queensland, Australia, maliciously released over 264,000 gallons of raw sewage, which flooded nearby rivers and parks [12]. In 2006, a foreign attacker installed malicious software remotely on systems at a water filtering plant in Harrisburg, Pennsylvania, negatively affecting the operations of the plant [12]. Because those systems are similar to dam systems, the same forms of attack can also target dam systems.

Supervisory Control and Data Acquisition (SCADA) systems are a major part of many industrial control systems. What makes them critical is the fact

that they are centralized and communicate with large-scale system components that typically operate on entire sites [9]. Components in most SCADA systems communicate with each other using the Modbus protocol, which offers no protections against denial-of-service (DoS) attacks and malicious data modification. Moreover, today's SCADA systems are based on open standards, so attackers can easily learn how they work.

To study the effects of specific attack behaviors on a water control system, we developed a case study of a river zonal SCADA dispatcher. It demonstrates that different system configurations can play a role in protecting a system, and that attack targets vary depending on the type of attacker. The main contribution of this work is an analysis of different attack scenarios and attacker types through stochastic modeling and quantitative metrics.

In Section 2, prior work related to the security evaluation of river zonal systems is discussed. Section 3 details the configurations that we considered. Section 4 introduces the modeling formalism, as well as the models used to analyze different attack scenarios and attacker types. Section 5 describes the experimental setup to carry out the analysis; Section 6 discusses the experimental results and analysis. Finally, Section 7 concludes the paper with a final discussion, including possible future improvements to this work.

2 Related Work

The authors of [1] study common attacks on SCADA control systems, such as command injection, data injection, and denial of service attacks. They used actual systems in a physical lab, such as a water storage tank control system, to study the effects of these attacks on the systems' operation, and designed an anomaly detection system that analyzes whether the data is normal or abnormal. Our approach is different in that it generalizes the anomalies by modeling attacks and intrusion detections and investigates how long it takes for an attack to unfold.

Common types of attacks on water control systems are provided in [6]. The authors grouped the attacks into four different classes: reconnaissance, response and measurement injection, command injection, and denial of service. *Reconnaissance attacks* gather control system information to help attackers locate hosts to attack, and *injection attacks* involve corruption of responses to make the system function abnormally. Reconnaissance and injection attacks were used in our study.

In [13], attacks on SCADA systems were described and classified by possible attack targets. Classifications include access control, memory protection, and protocol failure due to bugs. One interesting type of attack described in [13] is an Open Platform Communications (OPC) attack, in which an attacker compromises an HMI and then executes arbitrary code or attacks other servers. We used that attack to model how devices can be controlled by outsiders without the use of an HMI.

The authors of [7] described a river zonal dispatcher system, a SCADA traffic monitoring system for detecting potential intrusions, and a software agent model for the intrusion detection system for forensic analysis. We also considered the role of intrusion detection systems in the dispatcher, but we simply used the

concept to analyze quantitative timings so that we could determine how they play a role in inhibiting attacks.

3 Background

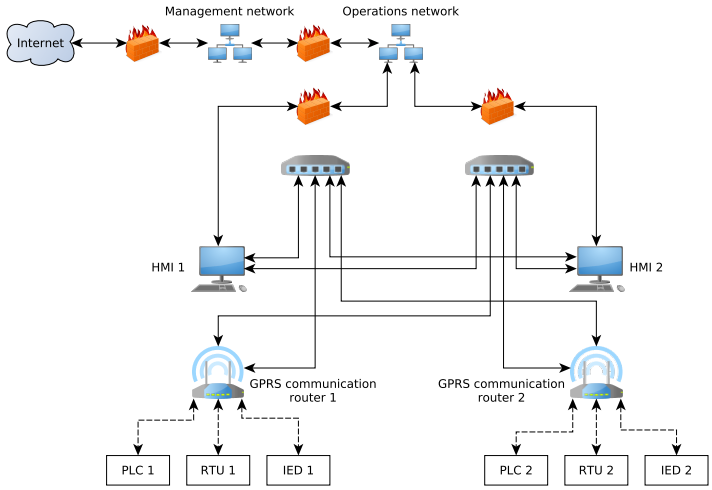

Fig. 1: River zonal SCADA dispatcher architecture without isolation.

The system we examined involves a river zonal dispatcher described in [7]. A simplification of its architecture is shown in Figure 1. The architecture consists mainly of four networks: a management LAN, an operations LAN, a supervisory LAN, and a bus network composed of on-site devices such as remote terminal units (RTUs), programmable logic controllers (PLCs), and intelligent electronic devices (IEDs). The management LAN consists of servers responsible for water resource management, and the operations LAN consists of servers collectively responsible for forecasting future outcomes given historical data. Each of the management LAN, operations LAN, and supervisory LANs are protected by an intrusion detection system (IDS) to detect possible intrusions. A supervisory LAN and a bus network together make up the meat of the system, and for this particular case study, there are two groups of these two types of networks. Each supervisory LAN consists of a human-machine interface (HMI), and each bus network consists of on-site devices. Each supervisory LAN is connected to its corresponding bus network through a General Packet Radio Service (GPRS) communication router, which is used for wireless communication between HMIs and devices that use Modbus, a widely used protocol for industrial communication systems [5]. In the Modbus protocol, an HMI sends special unencrypted packets containing a function code with parameters specific to the function of a device, and then the device typically replies back by echoing the original function code and returning the output data. Typically, communication

is done through an Open Platform Communications (OPC) server that converts outgoing HMI commands to the proper Modbus commands, and converts incoming replies back to a format that the HMI can recognize. Sensor networks are often used in SCADA systems to collect and use information such as hydrometric, pluviometric, and meteorological data, so the main purpose of the HMI is to use information from the sensors to decide how to control a particular device.

The first supervisory LAN controls one subriver basin, and the second supervisory LAN controls the other subriver basin; however, the first subriver basin is assumed to cascade into the second one. Therefore, the basin corresponding to the first supervisory LAN is referred to as the *upper* subriver basin, and the basin corresponding to the second supervisory LAN is referred to as the *lower* subriver basin.

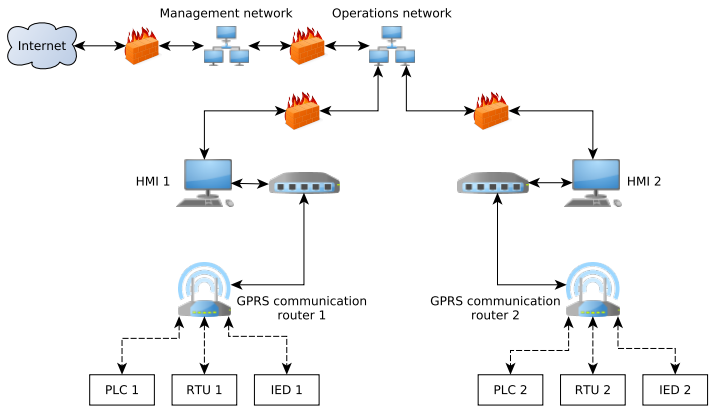

Fig. 2: River zonal SCADA dispatcher architecture with isolation.

Some assumptions were made: 1) the most critical part of the architecture is the SCADA system and is therefore the most desirable place for an adversary to target, which means that any given attack path from the Internet to the SCADA system is just for the purpose of gaining access to the SCADA system, and 2) the SCADA server and HMI reside on a single machine. [7] assumes that the SCADA systems are accessible from each other, but one potential way of minimizing the impact of attacks is by isolating the different SCADA systems in the dispatcher system, as illustrated in Figure 2. In this case, the operator of an HMI cannot use one HMI to access another HMI, and a one-to-one correspondence exists between HMIs and GPRS communication routers, so the operator can access devices through one router only.

4 Methodology

4.1 Attacks

The attack behavior in the system was modeled in the ADVISE formalism [2, 3], which is a stochastic model that assumes the perspective of an adversary

who selects an optimal attack based on its level of attractiveness compared to all other attacks available to him or her. The ADVISE formalism consists of an attack execution graph that contains a collection of attack steps. Attack step preconditions dictate whether it is possible at a certain time for an attack step to be executed, and outcomes make up the effects of the attack step. Access, skill, knowledge, and goal state variables are connected to attack steps in the attack execution graph, signifying that the connected state variables are used in the preconditions (arcs from state variable to attack step) or effects (arcs from attack step to state variable) of the attack step. In an attack execution graph, yellow rectangles represent attack steps, and red squares, green circles, blue triangles, and orange ovals represent access, knowledge, skills, and goals, respectively.

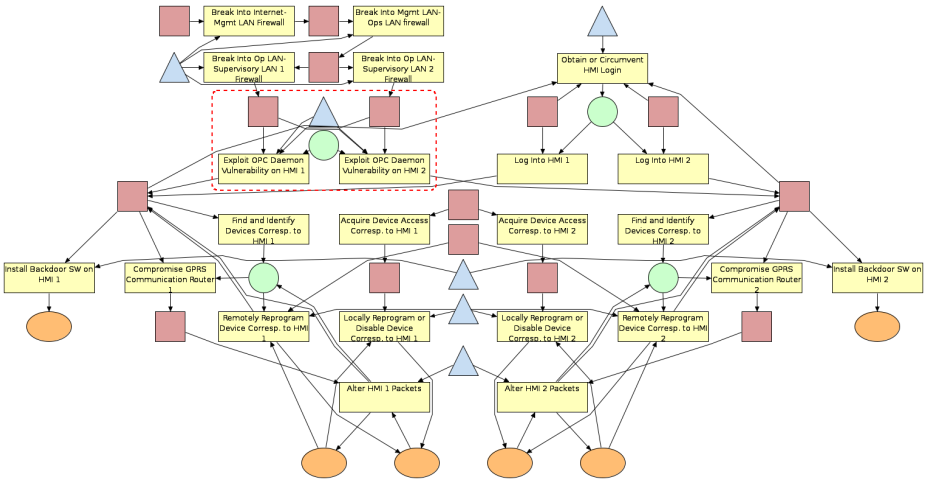

Fig. 3: ADVISE model of the river zonal SCADA dispatcher.

The ADVISE model for the river zonal dispatcher is shown in Figure 3. The section of the model in the dashed box exhibits different behavior depending on whether the supervisory LANs are isolated. The corresponding two cases are shown in Figure 4. Specifically, if the supervisory LANs are isolated, then exploitation of the OPC daemon vulnerability is not possible on supervisory LAN 1 through supervisory LAN 2 and vice versa. This means that the edge from *SupervisoryLAN1Access* to the HMI 2 exploitation step and the edge from *SupervisoryLAN2Access* to the HMI 1 exploitation step, seen in Figure 4a, have no effect on enabling those attack steps, so they effectively disappear, resulting in Figure 4b. However, if the network is not isolated (see Figure 4a), then exploitation of the OPC daemon vulnerability is possible on supervisory LAN 1 through supervisory LAN 2 and vice versa.

In order to attack the SCADA system, an outside attacker must first gain access to the management LAN, operations LAN, and any of the supervisory

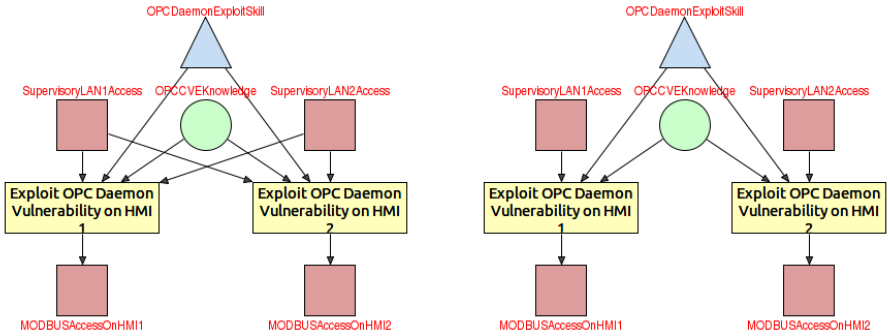

(a) HMI attacks without isolation. (b) HMI attacks with isolation.

Fig. 4: Detail of dashed section in Figure 3.

LANs containing the HMI. Those attacks correspond to attacking the firewalls that lie between the networks; those attacks are represented as four attack steps at the top left of Figure 3. Access to the supervisory LAN is one requirement for gaining control of the SCADA system. The ability to send commands to the on-site devices is another. The attacker may already have login access to an HMI (if he is an operator), or he can simply steal the password (if he is some other type of insider). If the attacker is an outsider, he can exploit the OPC server and issue commands to devices from there. In the case of non-isolated SCADA systems, gaining access to any of the HMIs grants the attacker access to all devices in the dispatcher, and in the case of isolated SCADA systems, gaining access to a single HMI grants the attacker access to devices through just one of the routers. Once the adversary has control of the HMI, he or she can attack the system in four possible ways: 1) install backdoor software on the HMI; 2) compromise the GPRS communication router that allows the HMI to interact with the devices, and then maliciously alter packets going through the router; 3) remotely reprogram the devices via the HMI so that they behave maliciously; or 4) directly reprogram the devices at the corresponding bricks-and-mortar facility so that they behave maliciously. In all, barriers to the attack include the need to defeat firewalls that connect pairs of networks together, to log into HMIs, or to compromise OPC servers, and to slip past IDSes that are placed on every LAN to detect possible intrusions.

Packets going through the router can be maliciously altered through injection attacks, as described in [6]. A more sophisticated attack, on the other hand, involves reprogramming of the devices at the facilities. If a device is equipped with function code 126 access [8], which allows a device to be remotely reprogrammed, then an attacker can simply use the HMI to find and identify the devices and remotely alter their behavior. If the attacker is an insider and has access to the bricks-and-mortar facility that houses a device, all he or she needs to do is enter the facility and directly reprogram the device. This type of attack satisfies the goal of compromising the device at that particular subriver facility.

One important assumption regarding these attack models is that no defense model is present, so when an attack goal is met, it remains met for the re-

mainder of time. However, the model considers the effects of IDSes, which affect attack steps involving the management LAN-operations LAN firewall, operations LAN-supervisory LAN firewalls, OPC daemon exploitation, HMI backdoor software installation, GPRS communication router compromise, router injection, and reprogramming of devices. In these cases, all accesses upon which the attack depends are restored to the initial state when those attacks are detected.

4.2 Response

A response model complements the ADVISE model and restores the state of the system after an attack is carried out. The response behavior in the system was modeled with the Stochastic Activity Network (SAN) formalism [4], which is an extension of Petri nets.

Whenever backdoor software is installed on any of the HMIs, it must first be detected, and then the actual repair process of uninstalling the backdoor takes place. The repair process restores the initial access that the adversary had; i.e., insider attackers still have access to the HMIs after the repair process completes, but outside attackers lose that access. Also, the repair process forces the attacker to re-achieve the goal of installing the backdoor software when he or she has the chance. Whenever the system is compromised via a device or router, someone must recognize that the system is operating abnormally, and then the actual repair process of restoring the functionality of the router or device takes place. Just like the repair process of uninstalling backdoors, the device and router functionality repair processes restore the initial access that the adversary had, and they force the attacker to re-achieve the goals of compromising the system via the device or router.

5 Experiment

To understand the different behaviors of each adversary and the various conditions that can increase the difficulty of an attack, simulations of the models were executed using different types of adversaries with and without IDSes present in the system. The goal of these experiments was to understand the types of scenarios that can negatively impact the security of the system and to determine the best practices for protecting it. Five types of adversaries were considered: 1) a foreign government, 2) a lone hacker, 3) a hostile organization, 4) an insider engineer, and 5) an insider operator. A foreign government is primarily concerned with installing backdoors on the HMIs and cares little about costs. A hacker is interested in most of the possible goals and is highly skilled, but must consider a balance of concern regarding cost, payoff, and detection. A hostile organization is also highly skilled, but is interested only in compromising the supervisory LANs and is mostly seeking the best payoff. The insider engineer is interested in all goals, but is poorly skilled in attacks, while the insider operator has access to many parts of the system already, is highly skilled, and is primarily concerned with reprogramming the devices.

There were a total of 20 cases for each simulation, as there are five different types of adversaries, supervisory LANs may be isolated or non-isolated, and IDSes may be present or not present. For each case, the percentage of the simulated time that an attacker has control of an HMI, router, or device was studied,

as well as the percentage of time the adversary takes to attack the system before reaching his goal. The systems were simulated for up to 8,760 hours (i.e., one year). The simulations were run for a minimum of 1,000 iterations and continued to run until either a 90% confidence interval for all measurements or a maximum of 10,000 iterations was achieved.

6 Results and Analysis

In the results that follow, a zero cost and detection probability of 0.95 were specified for the do-nothing step so that the attacker does not give up too easily when trying to accomplishing attack goals [2, 3].

6.1 Control of Device

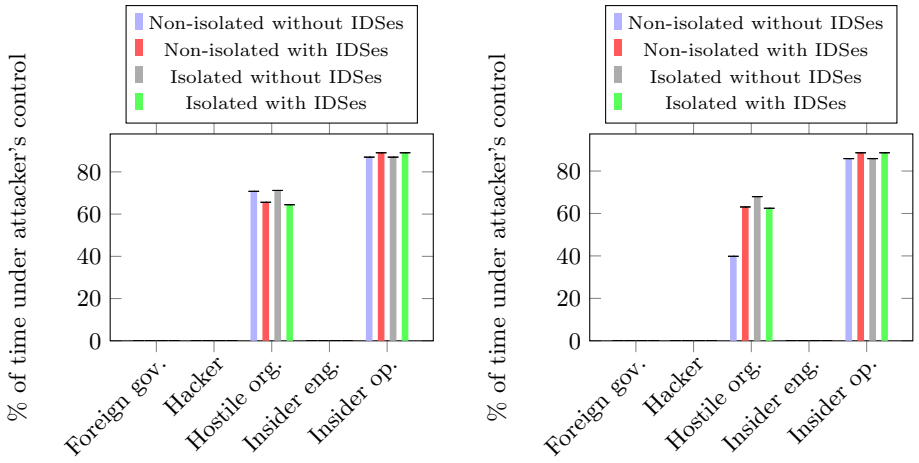

(a) Supervisory LAN 1.

(b) Supervisory LAN 2.

Fig. 5: Average percentages of time in which the attacker has control of an on-site device on a particular supervisory network.

Figure 5 shows the percentages of time that the subsystems were compromised via devices with respect to different attackers. As expected, the foreign government did not see gaining control of the system through devices as the most attractive goal. For the most part, the hacker also did not see any value in gaining control of the system through devices. However, we did observe one interesting result, which happened so rarely in the simulation that it can barely be seen in the figure: the attacker was able to penetrate the devices only when IDSes were enabled and indeed chose to do so, even though this attack was not in his best interests cost-wise, detection-wise, and payoff-wise. Specifically, at one point in the simulation, after finding the devices to compromise, rather than gain control of the system by launching router attacks or simply do nothing, the hacker saw that reprogramming the devices remotely was the most attractive course of action to take. The hostile organization proved to be successful

in reprogramming all the SCADA devices in the system. It reprogrammed the devices corresponding to HMI 2 far less often than the devices corresponding to HMI 1 in a non-isolated system without IDSes, but reprogrammed the devices with nearly equal interest in other scenarios. The insider engineer saw a large payoff in controlling the system through the devices, but because of low skill proficiency and cost considerations, he did not see it as the most attractive goal. The insider operator, on the other hand, has high attack proficiency for nearly everything (other than firewall attacks) and also had direct access to the facilities housing the devices, so he was able to gain control of the system without running into the obstacles faced by all other types of attackers.

IDSes proved to be not very effective in stopping device reprogramming attacks. Moreover, IDSes only helped the attacker stay more focused on goals with large payoffs. This is especially apparent when the hostile organization attacked the non-isolated networks in subriver system 2; the percentage of time in control jumped by 23% when IDSes were added. The lack of IDS effectiveness is also particularly obvious in the case of the insider operator, for which the percentage of time in control jumped by 3% when IDSes were added. This phenomenon results from the moderately high tolerance of detection by the hostile organization and an even higher tolerance of detection by the insider operator (i.e., their detection weights were 0.2 and 0.1, respectively). They were also highly skilled in attacking HMI components (with attack proficiencies of 0.7 or higher), which meant that it took very little time for them to bypass IDS protections in the supervisory LAN to gain further access.

The higher payoffs of attack goals for the upper subriver system did not make much of a difference under any attacker, since, for the most part, the attackers were constantly attacking the system until all possible goals were achieved.

6.2 Control of Router

Figure 6 shows the percentages of time that the subsystems were compromised via routers with respect to different attackers. Just as before, the foreign government did not see gaining control of the system through router attacks as the most attractive goal. The insider engineer and insider operator did not see much value in leveraging control of the system through router attacks, either. The hacker saw more value in performing attacks through the router corresponding to HMI 1, but did not see the same value in performing attacks through the other router (corresponding to HMI 2), although he found such attacks on the HMI 2 router to be more attractive (than the do-nothing step) in a non-isolated system with IDSes present. The hostile organization did not see much value in gaining control of the system through router attacks, as it is more interested in device reprogramming attacks, although there were several exceptions. First, the presence of IDSes made the attacker consider other types of attacks, such as device reprogramming attacks. Second, in the case of non-isolated, unprotected networks, the hostile organization saw as much value in router attacks as in device reprogramming attacks.

Overall, IDSes helped minimize router attacks in subriver system 1, but they did not perform as well in minimizing router attacks in subriver system 2. Specifically, for the hacker attacking subriver system 1, the presence of IDSes reduced the percentage of time in control by 10% in the non-isolated case and by 40% in

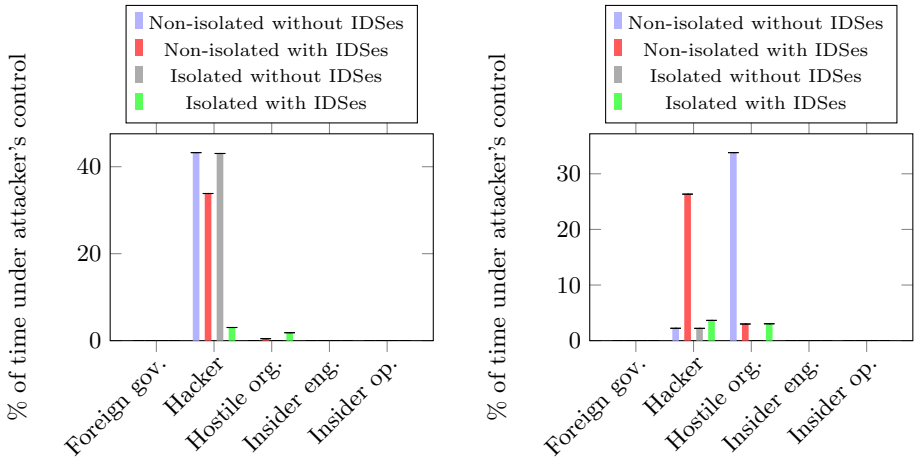

(a) Supervisory LAN 1.

(b) Supervisory LAN 2.

Fig. 6: Average percentages of time in which the attacker has control of a GPRS communication router on a particular supervisory network

the isolated case, which meant that isolation provided an extra layer of security in this scenario. On the other hand, when the same hacker attacked subriver system 2, the presence of IDSes increased the percentage of time in control by 24% in the non-isolated case and by less than 2% in the isolated case. Despite the increase in percentages, isolation helped provide an extra layer of security with a smaller increase in control of the system through router attacks.

6.3 Control of System via Backdoor Infection

Figure 7 shows the percentages of time that the subsystems were compromised via installation of HMI backdoors with respect to different attackers. Interestingly, every single type of attacker succeeded in installing backdoor software on the HMI in at least two different network configurations. However, this finding is unsurprising, because the goal of installing backdoor software is completely independent of the goals of controlling the system through devices that were directly altered by reprogramming or indirectly altered by router attacks, which meant the attackers had more leverage in installing backdoors on the HMIs.

The foreign government was interested in installing backdoors on both HMIs only when IDSes were not present. The hacker was more inclined to install backdoor software on the HMI in subriver system 1 over the HMI in subriver system 2. The hostile organization was also more inclined to install backdoor software on the HMI in subriver system 1, but only in cases where IDSes were present. Moreover, the hostile organization found that installing backdoor software on the HMI in subriver system 2 was one of the most attractive goals for an unprotected, non-isolated network configuration. Despite low proficiency in attack skills, the insider engineer was most successful in installing backdoor software on the HMIs, as he found it to be the least risky and most attractive goal. However, he was not interested in installing backdoor software on the HMI corresponding to subriver system 2 when IDSes were present. The insider operator was inter-

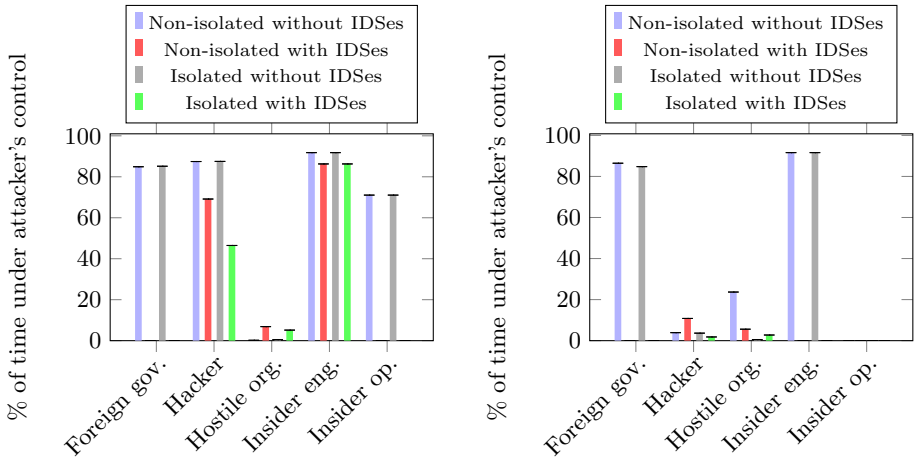

(a) Supervisory LAN 1.

(b) Supervisory LAN 2.

Fig. 7: Average percentages of time that an HMI on a particular network has backdoor software installed.

ested in installing backdoor software only on the HMI in subriver system 1 and only when IDSes were not present.

Overall, IDSes helped minimize backdoor software installations. They helped reduce system compromise via backdoor installations by anywhere from 2% to 40%; in some cases, they helped stop backdoor installations by the foreign government and insiders. However, there were a few exceptions. For example, a 7% increase was seen when IDSes were added in the case of backdoor software installation on the HMI in subriver system 2 by the hacker, and except when there were non-isolated HMIs, the hostile organization compromised the system through backdoor software installation for a slightly longer period of time when IDSes were present.

7 Conclusion

It was shown that network isolation and an IDS presence play a major role in the security of a system. We used ADVISE to study attacks on a river zonal dispatcher by investigating the effects of isolation and the presence of IDSes. This work is important because such systems send data to a national dispatcher that informs people of any danger, and the smallest difference in protection could have life or death consequences. If the system fails to function normally, warnings may cease, putting lives and property in serious peril. In many cases, IDSes help reduce the amount of time that a system remains in a compromised state, and isolation makes it more time-consuming for the attacker to explore attack paths. Mitigation of attacks through IDSes and isolation helps ensure that the national dispatcher can help save lives following catastrophes such as basin flooding.

The results indicated that attackers with certain abilities and focused goals are the most dangerous ones. For example, in the case of the insider operator, he not only had unrestricted access to the system, but also had a specific goal of attacking the devices. As a result, he was very successful in compromising

them. Further, the results suggest that security practitioners must account for as many types of adversaries as possible, since different types have different mindsets and different target goals. The best approach for security practitioners is to be proactive and keep their systems up to date.

Acknowledgments

The work described here was performed, in part, with funding from the Department of Homeland Security under contract HSHQDC-13-C-B0014, "Practical Metrics for Enterprise Security Engineering." The authors would also like to thank Jenny Applequist for her editorial efforts.

References

1. Gao, W., Morris, T., Reaves, B., Richey, D.: On SCADA control system command and response injection and intrusion detection. In: Proc. 2010 eCrime Researchers Summit (eCrime). pp. 1–9 (Oct 2010)
2. LeMay, E., Ford, M., Keefe, K., Sanders, W., Muehrcke, C.: Model-based security metrics using ADversary VIew Security Evaluation (ADVISE). In: Proc. 2011 Eighth International Conference on Quantitative Evaluation of Systems (QEST). pp. 191–200 (Sept 2011)
3. LeMay, E.: Adversary-Driven State-Based System Security Evaluation. Ph.D. thesis, University of Illinois at Urbana-Champaign, Urbana, IL (2011), http://www.perform.illinois.edu/Papers/USAN_papers/11LEM02.pdf
4. Meyer, J.F., Movaghar, A., Sanders, W.H.: Stochastic Activity Networks: Structure, behavior, and application. In: Proc. of the International Conf. on Timed Petri Nets. pp. 106–115. Torino, Italy (Jul 1985)
5. Modbus: Modbus application protocol specification v1.1b3 (Apr 2012), http://www.modbus.org/docs/Modbus_Application_Protocol_V1_1b3.pdf
6. Morris, T.H., Gao, W.: Industrial control system cyber attacks. In: Proc. 1st International Symposium on ICS & SCADA Cyber Security Research 2013. pp. 22–29. ICS-CSR 2013, BCS, UK (2013)
7. Stoian, I., Ignat, S., Capatina, D., Ghiran, O.: Security and intrusion detection on critical SCADA systems for water management. In: Proc. 2014 IEEE International Conference on Automation, Quality and Testing, Robotics. pp. 1–6 (May 2014)
8. Tenable Network Security, Inc.: Modicon Modbus/TCP programming function code access (2016), https://www.tenable.com/plugins/index.php?view=single&id=23819
9. U.S. Department of Homeland Security: Dams sector-specific plan: An annex to the national infrastructure protection plan (2010), http://www.dhs.gov/xlibrary/assets/nipp-ssp-dams-2010.pdf
10. U.S. Department of Homeland Security: Dams Sector (2015), http://www.dhs.gov/dams-sector
11. U.S. Department of Homeland Security: National infrastructure protection plan: Dams sector (Aug 2015), https://www.dhs.gov/xlibrary/assets/nipp_snapshot_dams.pdf
12. U.S. Environmental Protection Agency: Cyber security 101 for water utilities (Jul 2012), https://nepis.epa.gov/Exe/ZyPURL.cgi?Dockey=P100KL4T.TXT
13. Zhu, B., Joseph, A., Sastry, S.: A taxonomy of cyber attacks on SCADA systems. In: Proc. 2011 International Conference on Internet of Things and 4th International Conference on Cyber, Physical and Social Computing. pp. 380–388 (2011)