Cyberattacks on Primary Frequency Response Mechanisms in Power Grids

Article *in* Computer · November 2018

DOI: 10.1109/MC.2018.2876053

Varun Badrinath Krishna
University of Illinois, Urbana-Champaign

William Sanders
University of Illinois, Urbana-Champaign

Cyber-Attacks on Primary Frequency Response Mechanisms in Power Grids

Varun Badrinath Krishna, Ziping Wu, Vaidehi V. Ambardekar, Richard Macwan, and William H. Sanders

Information Trust Institute and Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign

E-mail: {varunbk, zw18, vaidehi3, rmacwan, whs}@illinois.edu

ABSTRACT. Primary frequency response is one of the most important mechanisms in today's power grids for making the grids resilient to faults, both malicious and non-malicious, that could lead to outages. In this article, we present the first study of cyber-attacks that could cause power outages by infiltrating operational technology (OT) networks in generation control systems to inhibit primary frequency response. Such attacks may have become possible only in recent years because generation controls, which were traditionally air-gapped, are becoming increasingly connected to OT networks. We evaluate those attacks and propose defense strategies.

Stuxnet was malware that was used to target centrifuges in a uranium enrichment facility in Iran in 2009, causing them to rotate beyond their safe limits and leading to physical damage. It is considered to be the first cyber-weapon that had an impact on the physical world, because it could connect to programmable logic controllers that caused damage to the centrifuges by operating them beyond their safety limits[1]. The second such cyber-weapon, called Crash Override, created an outage in Ukraine in 2015 through crafting of malicious commands that could cause grid relays to open, disconnecting loads from generators[2]. Those cyber-weapons have set a dangerous precedent and have motivated research on cyber-resilience for energy-delivery systems. In this article, we study a new attack vector that needs to be protected against in order to prevent cyber-induced power outages. Such outages could undermine critical defense infrastructure and much of the economy, and could place the health and safety of millions of people at risk.

Outages result when loads are disconnected from generators by protective equipment called *relays*, which isolate faults in a power grid. Faults typically cause the grid frequency to change drastically from its nominal value, which is 60 Hz in the U.S. If the frequency drops below a specified threshold, the relay disconnects the load so that a fault in one part of the grid does not have cascading effects on a different part of the grid. That disconnection is called *under-frequency load shedding (UFLS)*, and it can happen in the absence of attacks. Load shedding could also be caused by malware, such as Crash Override, that exploits vulnerabilities in specific relay software to cause the disconnection of loads. From the perspective of the electric transmission network, the load may represent a small town, and when UFLS disconnects that load, a regional outage could result. After UFLS, the total load connected to the generators is reduced, and the generators can be fully utilized to serve other loads in the grid (other towns) despite the disruption to the town in which the fault occurred.

There are frequency response mechanisms in power grids that provide resilience by restoring the grid frequency to prevent UFLS. In this article we present the first study of attacks that target the primary frequency response mechanism through malware similar to Stuxnet and Crash Override. The objective of the article is to provide generation operators and technology providers with an understanding of the risks of such an attack so that concrete steps can be taken to mitigate that risk.

A Brief Review of Frequency Response

There are three main mechanisms in today's power grids that make the grids resilient to faults that cause the frequency to drop. They are primary, secondary, and tertiary frequency response[3]. *Primary frequency response* (PFR) is a decentralized approach wherein generators monitor the grid frequency at their location and perform *droop control* to restore the frequency to its nominal value. Droop control essentially controls the rotation speed of the generator turbine by controlling the flow of water (in the case of hydro turbines) or steam (in the case of steam turbines). PFR takes effect immediately after the frequency has dropped below the nominal value. *Secondary frequency response* uses a centralized mechanism at the dispatch center to dispatch multiple generators from multiple areas of the grid to aid in frequency response. That mechanism is called *automatic generation control (AGC)*, and it takes 30 seconds to take effect after a fault. *Tertiary frequency response* uses reserve generators to compensate for any generation deficit and takes effect 10 to 30 minutes after the fault. This article is primarily concerned with PFR, whose responsibility is to ensure that the grid frequency is safely restored (to a value above the stipulated threshold for UFLS) before AGC takes effect. If PFR fails to do so, then UFLS would occur and cause regional outages *before secondary and tertiary frequency response have had the opportunity to restore the frequency*.

This article is the first to explore attacks on PFR. Turbine controllers that provide PFR have been analog and air-gapped from computer networks until recent years, and those air-gaps protected against cyber-attacks on PFR. However, those controls have been increasingly connected to distributed control systems (DCS) in generator operational technology (OT) networks. Therefore, this article is timely in that it provides prevention, detection, and response strategies for new vulnerabilities that may have emerged as a consequence of technology modernization and increased connectivity.

Under-Frequency Load Shedding Threshold

Different countries stipulate different thresholds for the frequency under which UFLS is triggered. In North America, the North American Electric Reliability Council (NERC) stipulates the threshold in the PRC-006-1 standard entitled "Automatic Under-frequency Load Shedding." According to that standard, after a major transient disturbance, UFLS will occur if the frequency falls below 0.575log(t) + 57.83 Hz after a period of t seconds. The logarithmic threshold accounts for the tendency of the frequency to recover slowly over a period of time as a result of control mechanisms that provide frequency response. Therefore, the threshold is initially lower (58 Hz for the first two seconds) and later higher (58.68 Hz after 30 seconds). If the frequency remains below 58.68 Hz at 30 seconds, UFLS will be triggered, potentially resulting in regional outages even before AGC can take effect.

System Description

In this section, we provide background on power grid analysis, generation controls, and network architectures in generation control centers.

Power Grid Analysis

The electric power transmission system comprises a multitude of generators and loads. The generators and loads are connected directly to buses, and the buses are connected by the transmission lines, transformers, and protective equipment. Broadly speaking, the power grid can be analyzed in two states: steady state and transient state. Frequency is not considered in steady state, because the system is assumed to be in a stable condition (i.e., no major frequency excursions). Transients, on the other hand, deal with unexpected failures that cause frequency excursions. Such failures can include short-circuit faults, sudden tripping or connection of large loads, and sudden opening of transmission lines, among others. In this article, we are interested only in failures that result in the disconnection of one or more generators.

Turbine Controls

Most electricity generators are powered by turbines. The turbines create a rotating magnetic field, which in turn produces alternating current. The turbines are driven by the movement of fluids, such as steam, hydro, and wind. The flow of those fluids can be controlled by manipulating valves (steam), dams (hydro), or blade angles (wind). For the same turbine movement resistance, allowing greater flow produces greater rotation speed and thus increases the frequency of the power generated. When a fault occurs in the grid, causing the disconnection of one or more generators, all connected generators' turbines experience increased resistance to movement, and the flow into the turbines needs to be increased in order to maintain the required frequency. As an analogy, if a chariot drawn by four horses were to lose one horse, the remaining three horses would feel a larger burden and need to exert greater effort in order to maintain the same chariot speed. This article is focused on steam turbine generators, which are by far the most common type. The steam in steam turbine generators is produced from heat from burning of fuels, such as coal and natural gas, or from nuclear fission.

To gain insight into the state of practice, we interviewed David R. Brown, a senior consulting engineer for turbomachinery and generator control applications at Schneider Electric. He has been working on turbine controls for over 40 years, and we learned from him that generation control operators can now use DCS to modify the maximum power output of each generator by setting a parameter, which we refer to as P_{MAX}. That modification is accomplished by regulating the steam flow into the turbine via valves to ensure that the power output does not exceed $\mathrm{P_{MAX}}$. The control of the flow of steam is key to providing frequency response, and that control has been analog and air-gapped until recent years. In recent years, according to Brown, the turbine controls have become increasingly automated with digital control and connectivity with DCS. The DCS provide connectivity between multiple generation units (each containing a turbine) and a centralized control center at the generation facility. That network architecture allows remote control of turbine control settings through human-machine interfaces (HMI). Companies that provide digital control equipment for generator turbines include Voith, Woodward, Schneider Electric, ABB, and Honeywell.

Communication Network Design in Generation Centers

Apart from setting UFLS and related standards, NERC also provides a set of critical infrastructure protection (CIP) standards for device manufacturers and operators. Critical infrastructure control centers that comply with those standards, particularly with the electronic security perimeter standard (CIP-005-5), have their information technology (IT) and OT networks segmented as illustrated in Figure 1.

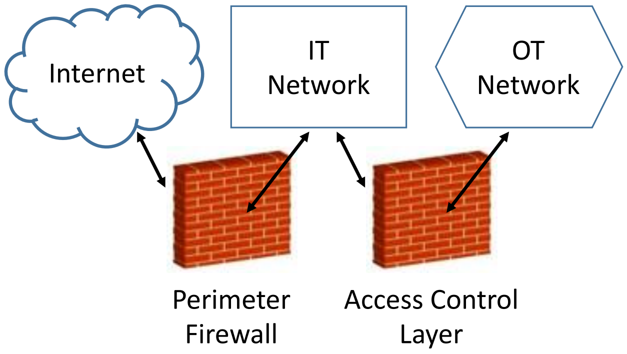

Figure 1: Partitioning of IT and OT Networks

As illustrated in Figure 1, the IT network has direct access to the Internet through a perimeter firewall. That firewall provides basic IT security, restricting inbound connections and optionally performing intrusion detection. An additional firewall (or a virtual private network with access controls) exists to separate the IT and OT networks, restricting access to specific users who have privileges to access the OT network. We assume that all generators comply with the NERC CIP standard. As a result, our baseline network security model reflects the current best practice.

Threat Model

In our threat model, we assume that the attacker has sufficient time and resources to plan and launch a sophisticated cyber-attack on power grid generators. Such an attacker could be a nation-state adversary, as was suspected for Stuxnet and Crash Override. The goal of the attacker is to create a regional outage by compromising PFR mechanisms in a subset of generators.

We assume that the attacker is unable to compromise the electric power transmission system operator facility (sometimes referred to as the *balancing authority*). Since those facilities have direct access to every sensitive piece of equipment in the power grid, gaining access to them would make it easy for an attacker to create an outage. As a result, we assume that great care has been taken to secure the system operator's physical and cyber territory.

We assume that in order to compromise generation facilities and cause an outage, the attacker would create an advanced persistent threat (APT) and take specific steps, as illustrated in Figure 2. The logic bomb ensures that the malware remains latent and is triggered at a future time, when the system is heavily loaded and is more likely to suffer an outage after a failure.

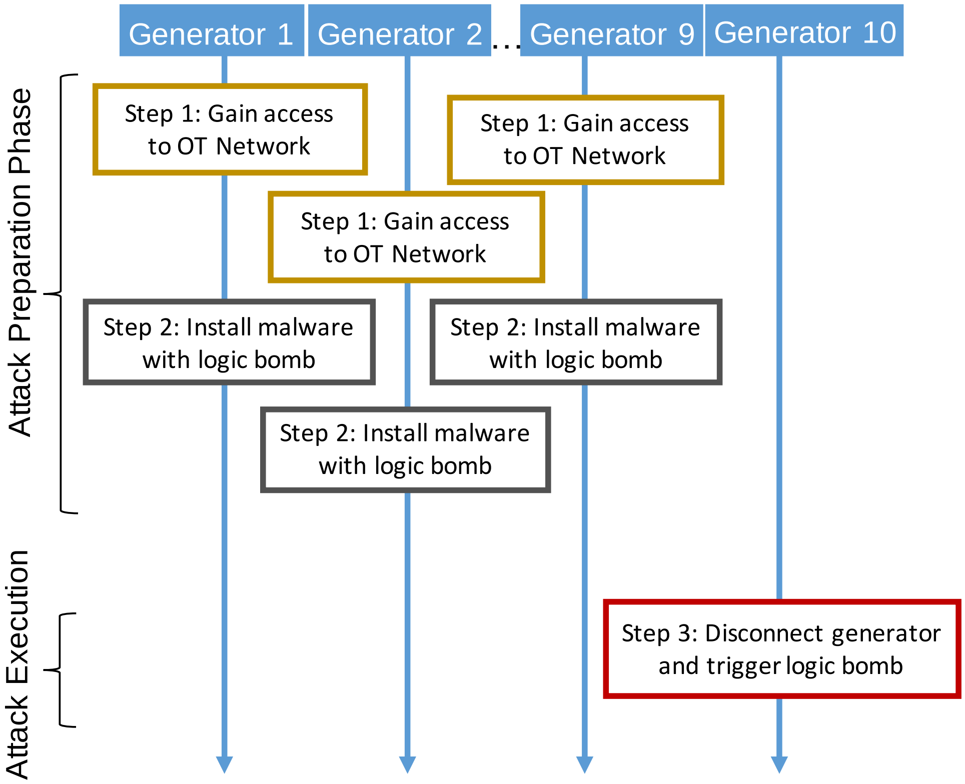
Generator 1
Generator 2
…
Generator 9
Generator 10
Attack Preparation Phase
Step 1: Gain access to OT Network
Step 1: Gain access to OT Network
Step 1: Gain access to OT Network
Step 2: Install malware with logic bomb
Step 2: Install malware with logic bomb
Step 2: Install malware with logic bomb
Attack Execution
Step 3: Disconnect generator and trigger logic bomb

Figure 2: Attack steps to cause an outage

Gaining Access to the OT Network

Although the NERC CIP standards are detailed and well thought-out, they make no mention of social engineering attacks, which are used by over 80% of hackers, according to a 2016 survey[4]. Therefore, an attacker could gain access to the OT network by first entering the IT network through social engineering attacks, such as phishing, spear-phishing, and baiting. That is precisely how Crash Override is believed to have entered the OT networks in Ukraine[2]. In order to gain access to the OT network from the IT network, the malware could enter a system being run by personnel with privileged access to the OT network. Without that access, as illustrated in Figure 1, the malware would not be able to move laterally out of the IT network and would not cause any physical harm. Once the malware has gained access to the OT network, it can send commands through the DCS to control the generation units.

Disconnection Attack

Once the attacker has gained access, they could simply instruct the DCS to disconnect all the generator units through generation circuit breakers. In order for that attack to cause an outage, the attacker would need to coordinate the disconnection of enough generators to create a UFLS scenario. In effect, the attacker would control a botnet, and synchronize the disconnection of multiple generators by communicating with the bot in each generator. While that attack might be effective, it would require management and coordination on the attacker's part. Also, it would not be subtle, and would raise suspicion of criminal intent.

At the time of this writing, there are very few generation circuit breakers that are connected to DCS systems. Therefore, the breakers may not be accessible to attackers, so disconnection attacks may not be feasible.

PFR Restriction Attack

We discovered a different attack model, wherein the malware targets the PFR system by reducing P_{MAX} and effectively restricting PFR. It is essential that the malware enter multiple generators in order to trigger UFLS, as we will show in our simulation study. To circumvent detection, the malware could be loaded with a logic bomb that would cause it to remain dormant and reduce P_{MAX} only on a specified day and time, by which time the attacker could ensure that the required number of generators have been compromised. The day and time could be chosen based on expected load conditions, which we will discuss in our simulation study.

After reducing P_{MAX}, the malware could also compromise the inputs to the HMI at the generation control center to mask the attack and make it appear as though P_{MAX} has not been modified. Ultimately, as illustrated in Figure 2, the attacker would need to cause a loss of generation after the logic bombs have been triggered in the malware. That could be accomplished by malware as described previously in the context of the disconnection attack. The disconnection attack would require the disconnection of multiple generators to be effective. The PFR restriction attack, on the other hand, could work even if only one generator was disconnected, and is therefore subtler than the disconnection attack. Being subtle, the attack could be misattributed to mismanagement of generators, or non-malicious faults, throwing off suspicion of criminal intent.

In comparison to the disconnection attack, the PFR restriction attack could take less effort for the attacker to implement because no coordination is required. As long as the malware has been installed in the attack preparation phase on all the target generators before the fault is induced, UFLS could be effected. The reason is that generators are inherently synchronized to the grid frequency; the attacker does not need bots that communicate to facilitate synchronization. After one generator has been disconnected in the attack execution phase, the change in inertia will be felt at all connected generators as the frequency declines. The other generators will then try to perform PFR by increasing their generation to compensate for the loss of generation in the system, but that increase will be limited by the value of P_{MAX}. In our simulation study, we will show that UFLS could result as long as a large number of generators have been compromised and their P_{MAX} settings have been set to values that are substantially lower than the default.

Simulation Study

We use the PowerWorld Simulator[5] because it can simulate both steady-state and transient scenarios in a power grid. Also, PowerWorld comes with widely used simulation models that have been validated, requiring us to make only minimal parameter modifications in order to demonstrate attacks on PFR. Since such attacks cause frequency excursions, we need to be able to analyze transient behavior, and PowerWorld provides sufficient fidelity to do so.

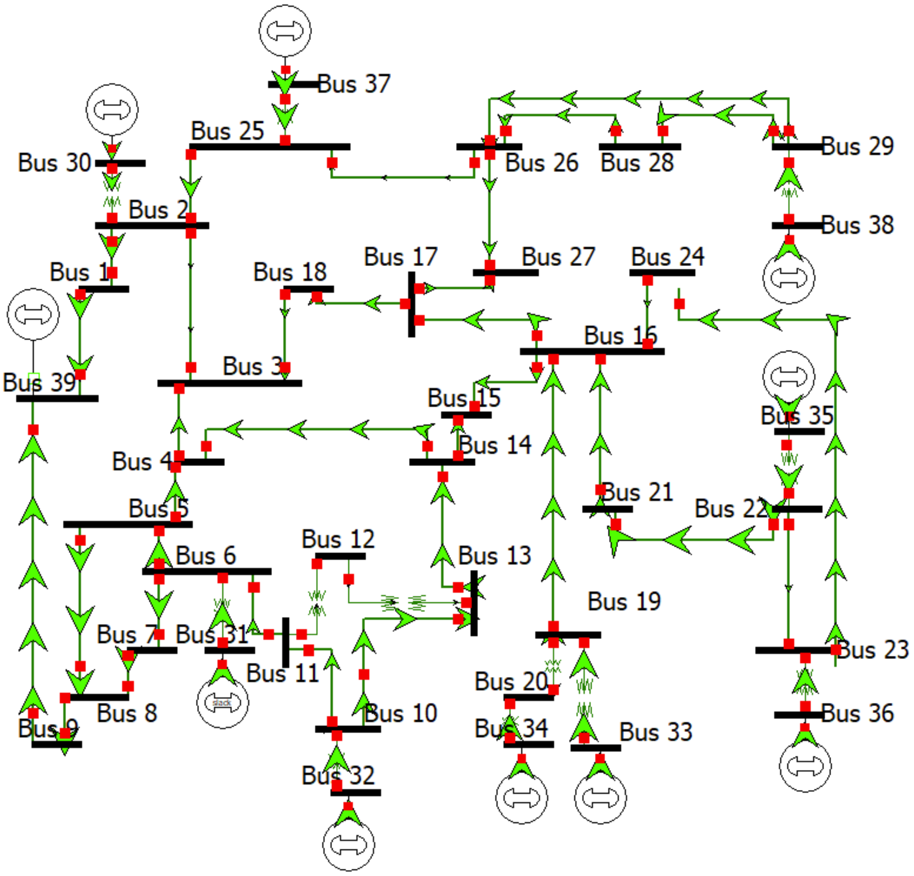

Figure 3: IEEE 10-Generator 39-bus New England Test System

Our power grid model is the standard IEEE 39-bus model of the New England grid, illustrated in Figure 3; it has 10 generators and 46 lines. We used the standard IEEEG1 governor control[6], in which the control parameter P_{MAX} limits the maximum generation in the transient state, as discussed previously in the section on turbine controls. We modified the generator capacities from the base system to illustrate the attack at different levels of the total load relative to the total generation capacity in the system. We refer to the ratio of the total load (including transmission losses) to the total generation capacity as the *relative load* of the system. The greater the relative load, the more stressed the system would be before the transient event. In our threat model, the transient event corresponds to the disconnection of a generator in Step 3 of the attack steps illustrated in Figure 2.

Three scenarios from our evaluation are illustrated in Figure 4(a). In all three scenarios, the default model was modified such that the relative load was 90%. We simulated a transient event in which the largest generator in the system (at Bus 39) was disconnected one second after the start of the simulation. As a result, the grid frequency started to decline from the nominal value of 60 Hz after one second. In the base case, PFR caused the frequency to stabilize at 59.77 Hz, which is a safe state above the UFLS threshold. That illustrates the inherent resilience of the grid to an attack that would cause the disconnection of the largest generator. As shown by the power flow arrows in Figure 3, the other generators compensated for the loss of that generator by providing additional power through PFR.

For the disconnection attack, two generators in addition to the one at Bus 39 needed to be simultaneously disconnected in order to cause UFLS. The PFR restriction attack illustrated in Figure 4(a) sets the P_{MAX} value to 90% of the generation capacity on 8 generators; no generator other than the one at Bus 39 was disconnected. In this experiment, the frequency fell below the UFLS threshold after 30 seconds, causing an outage right before AGC could take effect. Note that we illustrate the lower bound of the attacker's effort for all results presented in this section. Compromising more generators or reducing P_{MAX} to a greater extent would also cause UFLS, but that would require more effort from the attacker and possibly increase the risk of detection.

We now present a deeper exploration of the PFR restriction attack. We varied P_{MAX} as a percentage of each generator's capacity. In the base case (no attack), P_{MAX} was set to 120%, which is realistic because generators are allowed to exceed their capacity for a brief time period in order to enhance the system's resilience during a fault. If a generator were to operate beyond its rated capacity for too long, it could get damaged. For the attack cases, we modified P_{MAX}, varying it between 50% and 100% to determine settings that would result in UFLS. For each P_{MAX} and relative load setting, we used the simulator to obtain the minimum number of generators whose P_{MAX} values would need to be compromised in order to cause UFLS.

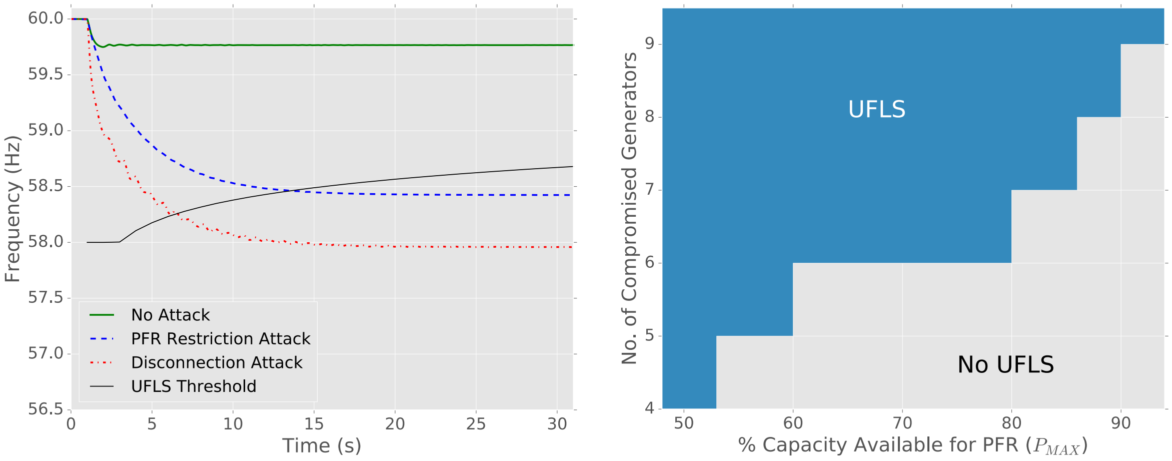

(a) Impact of Attacks on Grid Frequency

(b) Impact of PFR Attack Parameters

Figure 4: Summary of Results

The results are summarized in Figure 4(b) for a relative load set at 90%. The region corresponding to attack parameters for which there is a risk of UFLS is shaded in blue. To help the reader interpret the plot, we will describe two examples. First, restricting PFR by setting 60% ≤ P_{MAX} < 80% required that at least six generators be compromised to cause UFLS. Second, setting P_{MAX} = 90% restricted PFR less, so the system became more resilient, and the attack had to compromise all nine connected generators to cause UFLS. In general, fewer generators needed to be compromised when P_{MAX} was reduced to further restrict the generation capacity available for PFR.

Additional illustrations of other relative load levels can be found in the Ph.D. dissertation associated with this article[7]. As expected, greater relative load caused the UFLS region to become larger because the system was operating closer to its generation capacity and was unable to tolerate small reductions of P_{MAX}. The UFLS regions for smaller relative loads were fully contained in the UFLS regions for larger relative loads. That indicates that an attacker is more likely to induce UFLS if they were to define a logic bomb (in Step 2 of Figure 2) to activate PFR restriction and induce the fault (in Step 3 of Figure 2) when the system is expected to operate at high loading conditions.

Defense Strategies

We now discuss prevention, detection, and response for the attacks presented.

Prevention

We suggest that utilities not only comply with the NERC CIP standard, but also augment standard personnel training with additional training on social engineering attacks, such as phishing, spear-phishing, and baiting. Further, e-mail filters in the IT network can be configured to block suspicious e-mails containing attachments or links to unknown Web domains. That would reduce the risk of downloading malware onto the IT network. In addition, we propose further segmentation of the IT network to ensure that personnel with access privileges to the OT network use a different, lower-privileged account when logging into computers for the purpose of browsing the Web and checking e-mail. That would ensure that any malware that might have been accidentally downloaded onto those computers would not have the privileges it needs to communicate with devices on the OT network.

If the logic bomb that is used to trigger the malware depends on remote control, the malware will need to communicate with the controller at an external IP address. Restricting outbound connections to unknown IP addresses would help prevent such remote control.

For generation controls in particular, additional measures can be taken. P_{MAX} is not altered on a frequent basis, so it can be made read-only through digital interfaces. If it does ever need to be altered, the change can be done manually, as it was for decades until the recent advent of turbine control automation.

Detection

If an attacker were to modify P_{MAX} and also cause the HMI to show that no modification had taken place, it would be very difficult for the generation operator to detect the attack. We propose the use of an air-gapped measurement device such that there is no communication path through which the attacker can compromise that device. In particular, a tachometer can be installed to measure the number of rotations per minute, which is directly proportional to the grid frequency at the generator. Any drop in the frequency resulting from a malicious reduction of P_{MAX} will be reflected in the tachometer reading, and alerts can be sent to technical staff if the frequency drops below preconfigured thresholds.

Response

If the system is believed to have been compromised, the operator can respond by overriding all controls to known safe settings and falling back to a manual operation mode wherein the controls are air-gapped from the OT network. The connection to the OT network can be restored after the malware has been removed or quarantined.

Additional Threat Models

In this section, we briefly discuss additional threats to power grid resilience.

Attacks on AGC

Attacks on AGC are relevant only if UFLS is not triggered within 30 seconds. Even if an attacker were to compromise all the generators in one area of the power grid, AGC would allow for power to be fed into that area from other areas. AGC uses the area control error (ACE) to determine how much power needs to be delivered to the affected area (call it *Area 1*) from a healthy area (call it *Area 2*). An attacker could compromise the ACE reading to make it appear to the operator that there is no need for Area 2 to support Area 1. The ACE reading is a single point of failure, and compromising it would ensure that the generators in Area 2 will fail to receive the AGC signal from the operator to increase generation to compensate for the lost power in Area 1. PFR, on the other hand, is distributed, requiring the attacker to compromise multiple generators in order to create an outage. Because AGC is centralized, it is not only easier to attack, but also easier to defend, because security mechanisms can be focused on maintaining the integrity of the single point of failure. PFR, on the other hand, requires that multiple, possibly heterogeneous turbine controllers be secured in order to prevent attacks. Attacks on AGC have been discussed in the literature[8].

Electricity Theft

Although the primary motivation for stealing electricity is monetary gain, the theft of large amounts of electricity can undermine the resilience of the grid. That was demonstrated in 2012 by the world's largest-ever outage, during which 670 million people in India lost electric power largely because of consumption via theft that was unaccounted for in generation scheduling[9]. Theft detection using smart meters and data science has been studied in the literature[10, 11].

The dissertation[7] associated with this paper contains additional evaluations, simulation details, and references to related work.

Acknowledgments

This material is based upon work supported by the Department of Energy under Award Number DE-OE0000780. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof. We thank David R. Brown, Prof. Peter Sauer, Jeffery Dagle, and Jenny Applequist for their insights and feedback.

References

1. K. Zetter, "An Unprecedented Look at Stuxnet, the World's First Digital Weapon," *Wired*, Mar. 3, 2014. [Online]. Available: https://www.wired.com/2014/11/countdown-to-zero-day-stuxnet/ .

2. A. Greenberg, "'Crash Override': The Malware That Took Down a Power Grid," *Wired*, Jun. 12, 2017. [Online]. Available: https://www.wired.com/story/crash-override-malware/.

3. Z. Wu, W. Gao, T. Gao, et al., "State-of-the-Art Review on Frequency Response of Wind Power Plants in Power Systems." *Journal of Modern Power Systems and Clean Energy*, vol. 16, no. 1, pp. 1-16. 2018.

4. J. Goldman, “Fully 84 Percent of Hackers Leverage Social Engineering in Cyber Attacks,” *eSecurity Planet*, Feb. 28, 2017. [Online]. Available: https://www.esecurityplanet.com/hackers/fully-84-percent-of-hackers-leverage-social-engineering-in-attacks.html.

5. PowerWorld Corporation, “PowerWorld Simulator,” [Online]. Available: https://www.powerworld.com/

6. IEEE, “IEEEG1 Governor Model,” [Online]. Available: https://www.powerworld.com/WebHelp/Content/TransientModels_HTML/Governor\%20IEEEG1\%20and\%20IEEEG1_GE.htm.

7. V. Badrinath Krishna, “Data-Driven Methods to Improve Resource Utilization, Fraud Detection, and Cyber-Resilience in Smart Grids,” Ph.D. dissertation, University of Illinois at Urbana-Champaign, 2018, to appear.

8. R. Tan, H. H. Nguyen, E. Y. S. Foo, D. K. Y. Yau, Z. Kalbarczyk, R. K. Iyer, and H. B. Gooi, “Modeling and Mitigating Impact of False Data Injection Attacks on Automatic Generation Control,” *IEEE Transactions on Information Forensics and Security,* vol. 12, no. 7, pp. 1609-1624, 2017.

9. B. Jairaj, “India’s Blackouts Highlight Need for Electricity Governance Reform,” World Resources Institute, Aug. 13, 2012. [Online]. Available: http://www.wri.org/blog/2012/08/india%E2%80%99s-blackouts-highlight-need-electricity-governance-reform.

10. V. Badrinath Krishna, K. Lee, G. A. Weaver, R. K. Iyer, and W. H. Sanders, “F-DETA: A Framework for Detecting Electricity Theft Attacks in Smart Grids,” in *Proc. 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, Toulouse, France, 2016, pp. 407-418.

11. V. Badrinath Krishna, C. A. Gunter, and W. H. Sanders, “Evaluating Detectors on Optimal Attack Vectors that Enable Electricity Theft and DER Fraud,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 12, no. 4, pp. 790-805, Aug. 2018.

Varun Badrinath Krishna is a Ph.D. candidate in the Department of Electrical and Computer Engineering and a Research Assistant in the Information Trust Institute at the University of Illinois at Urbana-Champaign. His research interests are in making critical infrastructures more efficient and secure using data-driven techniques. He is a Student Member of the IEEE.

Dr. Ziping Wu is a Power Research Engineer in the Information Trust Institute at the University of Illinois at Urbana-Champaign. His research interests include power system operation and control, renewable power generation, cyber-physical security, PMU applications, and microgrids. He is a Senior Member of the IEEE.

Vaidehi Ambardekar is an undergraduate research assistant in the Department of Electrical and Computer Engineering at the University of Illinois at Urbana-Champaign.

Mr. Richard Macwan is a Research Engineer-Power in the Information Trust Institute at the University of Illinois at Urbana-Champaign. His current research is mainly aimed at developing realistic cyber-physical testbed environments for implementation and assessment of algorithms for enhancing the cyber-security of power grids. He is a Member of the IEEE.

Prof. William H. Sanders is a Donald Biggar Willett Professor in Engineering and the Head of the Department of Electrical and Computer Engineering at the University of Illinois at Urbana-Champaign. He is a Fellow of the IEEE, the ACM, and the AAAS. His research interests are in assessment-driven design of trustworthy cyber infrastructures for societal-scale systems.