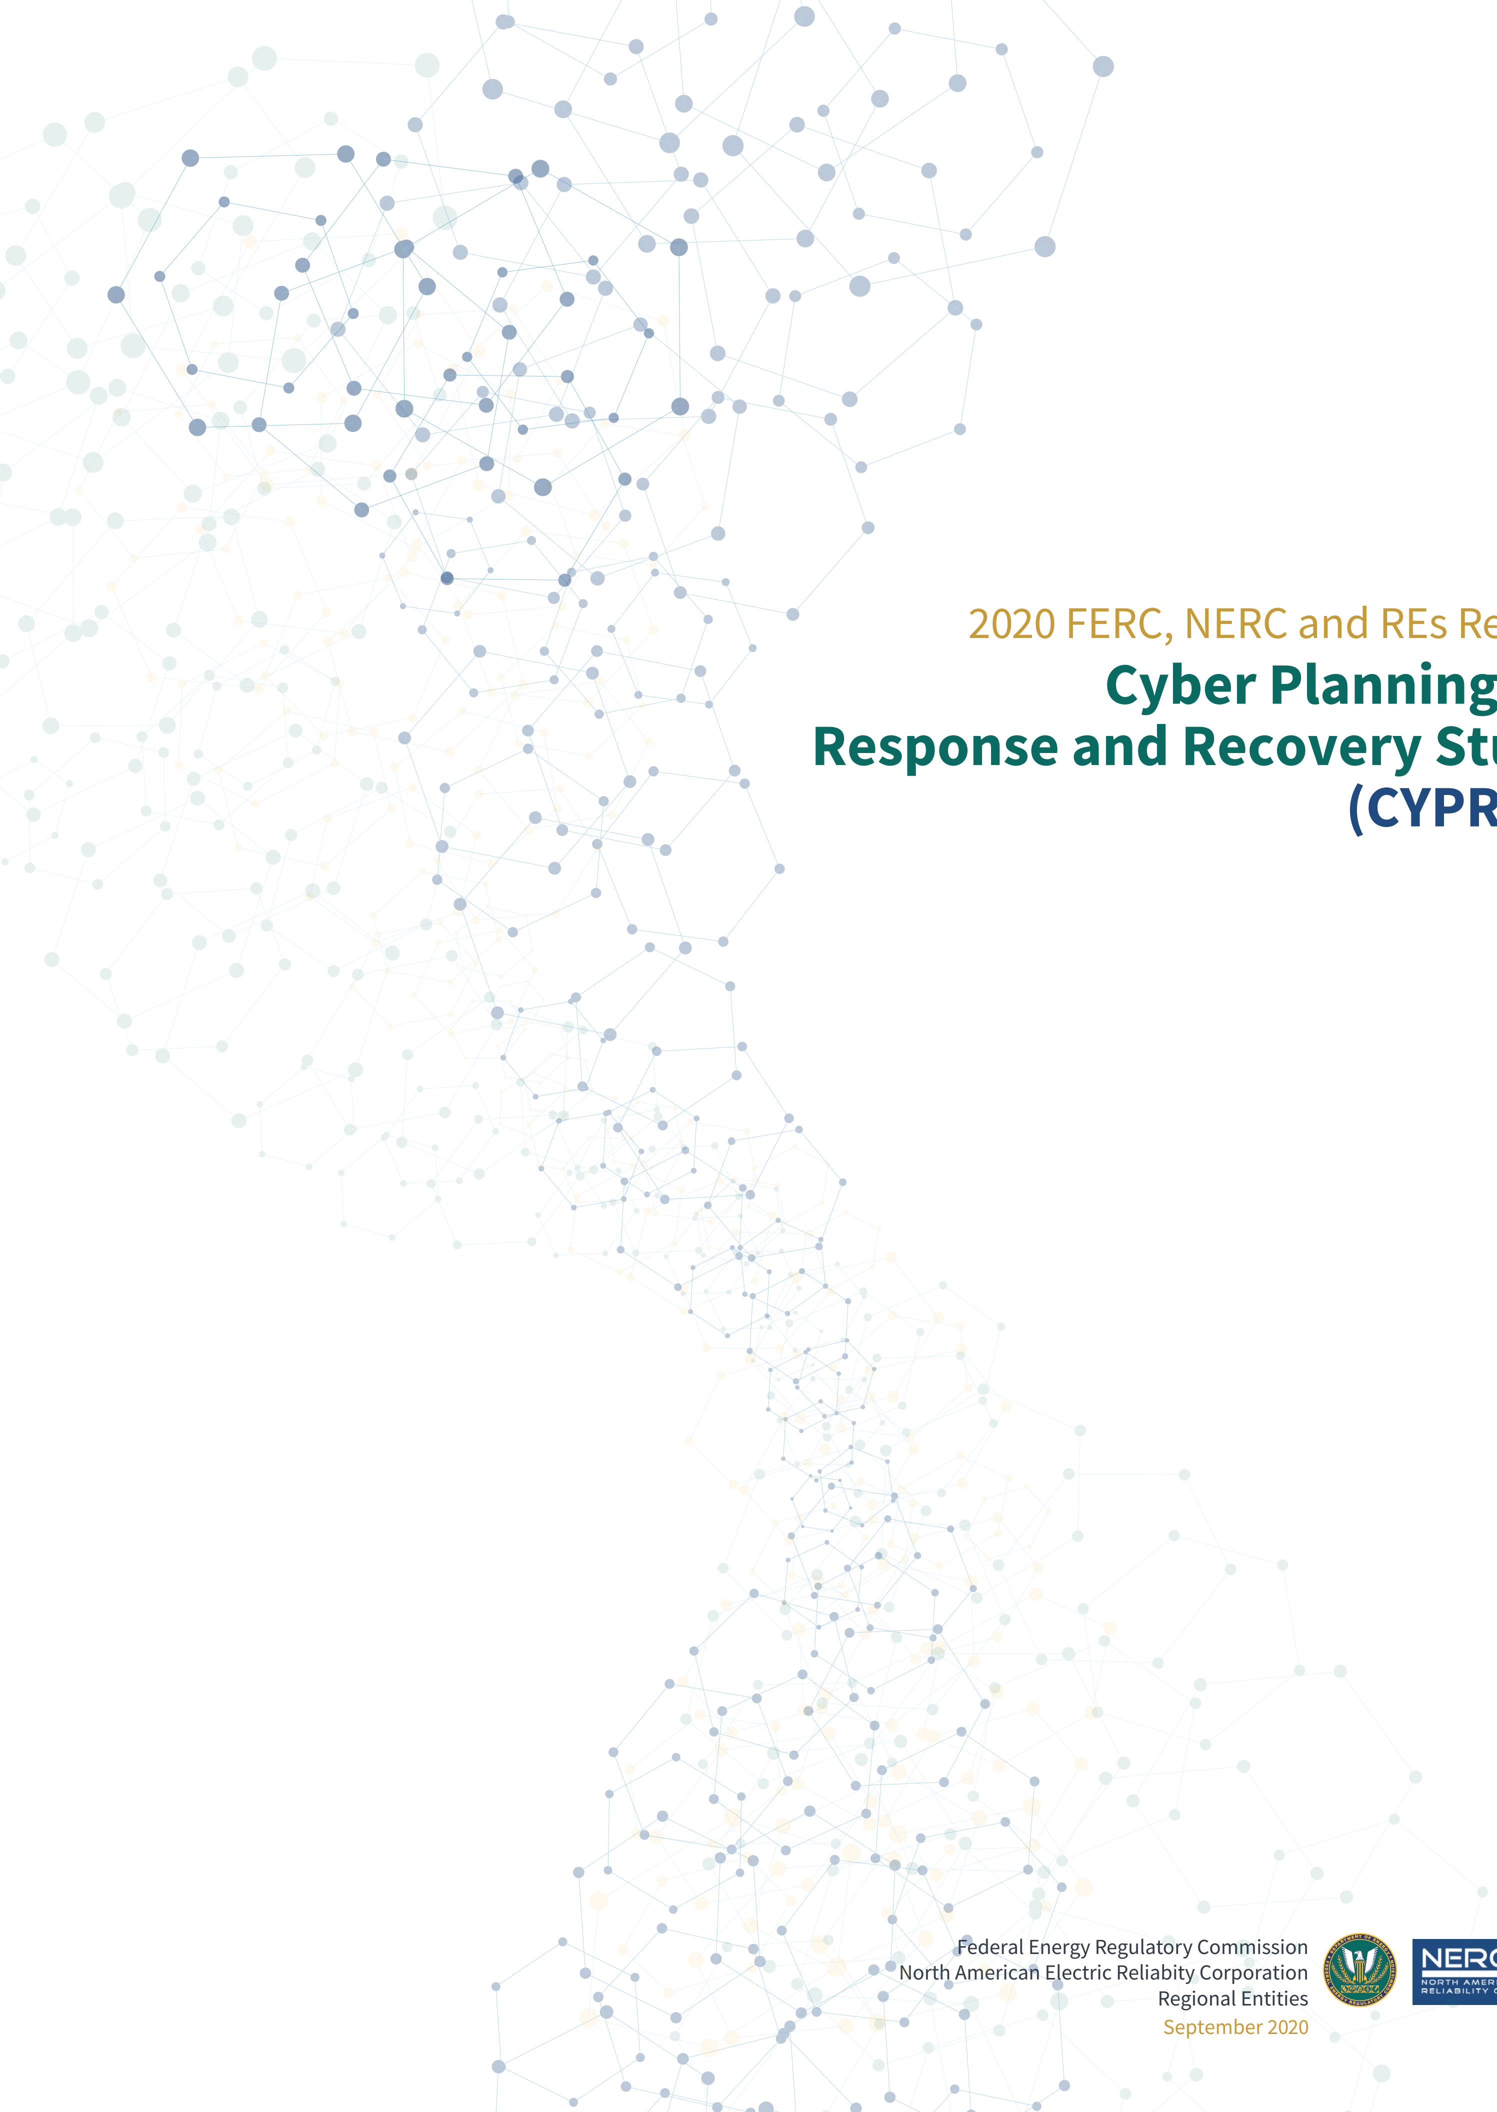

2020 FERC, NERC and REs Report

Cyber Planning for Response and Recovery Study (CYPRES)

Federal Energy Regulatory Commission
North American Electric Reliability Corporation
Regional Entities
September 2020

Cyber Planning for Response and Recovery Study

(CYPRES)

2020 FERC, NERC and REs Report

September 2020

FEDERAL ENERGY REGULATORY COMMISSION

Neil Chatterjee, Chairman
Richard Glick, Commissioner
James Danly, Commissioner

NORTH AMERICAN ELECTRIC RELIABILITY CORPORATION

Contents

Executive Summary

The Federal Energy Regulatory Commission (Commission) initiated a joint staff review, in partnership with the North American Electric Reliability Corporation (NERC) and its Regional Entities, to discuss their methods of incident response and recovery. This report was developed based on the team's observations.

Cyber threats pose a risk to electric utilities because they can impact operations and impose substantial costs. An incident response and recovery (IRR) plan describes how a utility responds to a cyber incident and includes phases and procedures. Establishing clear procedures for handling incidents is a complex undertaking and, though individualized to an organization's mission, size, structure, and functions, generally contain common elements: (1) they define their scope (to whom they apply, what do they cover, and under what circumstances); and (2) they define computer security events and incidents, staff roles and responsibilities, levels of authority for response (e.g., authority to disconnect equipment), reporting requirements, requirements and guidelines for external communications and information sharing, and procedures to evaluate performance.

The joint team consisting of Staff Members from FERC, NERC, and the Regional Entities conducted interviews with eight different entities varying in size and function to better understand their approach to IRR and to share the results with industry to enhance IRR planning generally. While there is no best method to develop, implement, or utilize plans, the goal of any IRR plan is to ensure the reliability of the Bulk Electric System (BES) in the event of a cyber security incident. While the IRR plans shared many similarities as to their development, implementation, and utilization, each entity's IRR plan had some differences in content and definitions for a cyber security event and incident. Additionally, some entities had separate IRR plans for their operational environments and business networked environments. Others utilized IRR plans that were very similar, with only slight variations specifically addressing the need for high availability of the operational networks.

Regardless of the approach, the entities followed a similar categorized framework identified in the National Institute of Standards and Technology (NIST) Special Publication 800-61 Revision 2. Using this framework, the joint team observed the following key take-aways.

Preparation:

- Effective IRR plans contain well-defined personnel roles, promote accountability, and, where appropriate, empower personnel to take action without unnecessary delays.
- Effective IRR plans leverage technology and automated tools while also recognizing the importance of human performance.
- Effective implementation of IRR plans requires well-trained personnel who are constantly updating their skills.
- Effective IRR plans incorporate lessons learned from past cyber security incidents or tests (drills, tabletop and operational exercises).

Detection and Analysis:

- Baselining is an effective resource utilization tool that allows personnel to detect significant deviations from normal operations.
- Flow-charts or decision trees are useful tools to determine quickly when a predefined risk threshold is reached, and a suspicious set of circumstances qualifies as an event.

Containment and Eradication:

- If an IRR plan containment strategy includes islanding operational networks (i.e. removing all external connections), there should be a thorough understanding of the potential impact of such a decision (e.g., through tests and training).
- IRR plans should consider the possibility that a containment strategy may trigger predefined destructive actions by the malware.
- Evidence collection and continued analysis are important to determine whether an event is an indicator of a larger compromise.
- IRR plans should consider the resource implications of incident responses of indeterminate length.

Post-Incident Activity:

- Effective IRR plans implement lessons-learned from previous incidents and simulated activities identifying clear shortfalls in the IRR plan.

I. Introduction

The joint staffs of the Federal Energy Regulatory Commission (Commission), North American Electric Reliability Corporation (NERC), and the Regional Entities (REs) have prepared this Cyber Planning for Response and Recovery Study (CYPRES) to assess the planning and readiness of electric utilities to respond to and recover from a cyber security incident. In September 2014, the Commission, NERC, and the REs initiated a joint staff review to assess electric utilities' plans for restoration and recovery of the bulk-power system following a widespread outage or blackout.[1] That effort culminated in the issuance of a joint Review of Restoration and Recovery Plans (Restoration and Recovery Report) in January 2016.[2] The Restoration and Recovery Report provided a comprehensive understanding of the electric utility industry's bulk-power system recovery and restoration planning, concluding that entities have system restoration plans that, for the most part, are thorough and highly detailed.

The Restoration and Recovery Report made several recommendations to improve system restoration and cyber incident response and recovery planning and readiness; for example, it recommended conducting exercises of response and recovery plans. Moreover, the Restoration and Recovery Report recommended conducting a study "to better understand plan improvements made by entities where testing or an actual cyber event revealed the need or opportunity for improvements" and to "examine and identify best practices with regard to the types of plan improvements made from entities' analyses of actual cyber events and/or testing."[3] In response to that recommendation, this report describes the methodologies implemented by entities in IRR plans.

II. Development Process for the Report

A joint team of subject matter experts from the Commission, NERC, and the REs (collectively, joint team), prepared this report using the National Institute of Standards and Technology (NIST) Special Publication 800-61 Revision 2, Computer Security Incident Handling Guide Recommendations, as a reference tool.[4] The NIST publication helps organizations mitigate the risks from computer security incidents by providing guidelines on how to respond to incidents effectively. The joint team conducted site visits to interview employees at electric

[1] The terms bulk-power system and bulk electric system are used interchangeably in this report and generally refer to facilities necessary for electric transmission but not local distribution.

[2] *Report on the FERC-NERC-Regional Entity Joint Review of Restoration and Recovery Plans*, January 2016, at http://www.ferc.gov/media/news-releases/2016/2016- 1/01-29-16.asp (Restoration and Recovery Report).

[3] *Id.* at 101-102.

[4] NIST Special Publication 800-61, Revision 2, Computer Security Incident Handling Guide at 21 (Figure 3-1. Incident Response Life Cycle) (NIST Special Publication 800-61), https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r2.pdf.

utilities[5] who oversee restoration and recovery planning. For completeness, the joint team selected entities and solicited volunteers with varying cyber infrastructure designs, significant BES responsibilities, and multiple NERC-registered functions.[6] To facilitate an open discussion of the participants' approaches to restoration and recovery planning, which entails initial detection of a suspicious event through the response and recovery process, the joint team agreed not to disclose entity-specific information outside of each review group. Accordingly, this report provides the results of the site visit observations without attribution to individual entities.

A. Interviews

The joint team conducted interviews with the entities' subject matter experts. To ensure consistency, the joint team conducted the interviews using the Cyber Kill Chain® intrusion process to guide discussions.[7] The joint team asked participants to explain how they would approach various phases of an intrusion in their incident and response recovery plans. Participants were asked about their corporate networks (sometimes referred to as an information technology (IT) network), as well as their operational technology (OT) networks, which control power systems. The joint team also asked participants how they leveraged NERC's Critical Infrastructure Protection (CIP) Reliability Standards, including Reliability Standard CIP-008-5 (Incident Reporting and Response Planning) in their IRR plans. This interview framework allowed the joint team to understand each participant's defensive capabilities in response to actions that an adversary might take as part of a cyber intrusion. The joint team then compiled a set of observations on the participants' defensive capabilities and the effectiveness of participants' response and recovery actions to a cyber intrusion based on the responses to interview questions and subsequent discussions among the joint team. The joint team also discussed logistics needed to support incident response, such as examining how employees who oversee and maintain various aspects of the plans ensure proper testing, appropriate staffing levels, funding, and training.

B. Background on Cyber Networks

To better appreciate the observations contained in this report, it is helpful to have a basic understanding of cyber network environments, including corporate IT systems and OT cyber

[5] The eight entities interviewed by the joint team, collectively, are Primary Compliance Contacts, Balancing Authority, Distribution Provider, Generator Owner, Generator Operator, Load-Serving Entity, Reliability Coordinator, Transmission Owner, and Transmission Operator.

[6] The NERC Functional Model defines the functions that must be performed to ensure the reliability of the bulk-power system. NERC, Reliability Functional Model, Version 5 https://www.nerc.com/pa/Stand/Functional%20Model%20Archive%201/Functional_Model_V5_Final_2009Dec1.pdf.

[7] The Cyber Kill Chain® process was developed by Lockheed Martin. https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html.

systems. Below, the joint team provides an overview of network design and system operations while also defining certain foundational terms and concepts.

Cyber Network Design

For the purposes of this report, cyber network design is defined as interconnected digital components such as computers, routers, switches, and firewalls. The design of each entity's cyber network is different due to various factors, such as the age of the system, personnel expertise, and the entity's size, mission and function. Nevertheless, cyber network systems have commonalities: they have a corporate IT network and an OT network. An IT network is used for day-to-day administrative tasks needed to run the utility.[8]   And in the utility context, an OT network is used to manage systems that provide services to customers, such as electric power generation controls and control room operations.[9]   For example, OT networks provide services for energy management systems, generation management systems, transmission systems and distribution systems. Entities implement these networks and systems to support reliability services provided by personnel at primary and backup control centers, data centers, generation facilities, and field facilities, such as transmission substations.

Corporate IT and OT networks have different levels of security, which are known as "trust-zones," because the data that resides and traverses these networks are used for different purposes. Corporate IT networks typically provide services to the user - the utility employee - such as web browsing, email, and social media access. These services allow the user to interact directly with the Internet. In contrast, OT networks typically restrict access to services that allow a user to directly interact with the Internet. The Internet, instead, is used as a pathway only to reach approved sites provided by a third party or to reach protected connections to industrial control systems that operate the BES. These trust-zones are a fundamental design feature for protecting data and systems with different security levels.

[8] While participant corporate IT networks were not the focus of this study, the joint team discussed these systems to understand the connectivity, interaction, data flow, and resource sharing (if any), between the participants' corporate IT and operational networks. In most cases, the participants explained that, while they deployed similar systems in the corporate IT and operational networks, the systems are discrete to those environments and do not share hardware or software capabilities. In some instances, participants explained that corporate IT systems, such as phone systems and Physical Access Control Systems, are configured to support both networks.

[9] All entities that participated in the study segment or separate their operational networks from the corporate IT networks. All participants described the operational networks as separate and distinct infrastructures, including LAN/WAN cabling, network switches and routers, servers and workstations, and storage systems. Participants using virtualized systems implemented those systems to share and leverage hardware and software resources within the operational network to create virtual hosts and systems for more effective system provisioning and recovery. No participants identified "mixed trust" virtualized systems between corporate IT and operational networks.

NERC's CIP Reliability Standards require enhanced trust-zones to protect the critical systems that operate the BES.[10]

Exercises

Exercises are opportunities for entities to test their systems and train their staff to ensure that their IRR plans appropriately mitigate risk and respond to a range of events. Participants can engage in exercises to test their response and recovery plans either within their respective regions or internally. The NERC Grid Security Exercise (GridEx) is an example of an exercise.[11] The joint team used the process and results from participants' exercises as a basis for discussion.

Cybersecurity Events and Incidents

In the cybersecurity context, entities use different definitions of "event" and "incident." This is because entities have different system and network designs and different approaches to responding to an "event" or "incident" depending on the risk to the reliability of their operations. For example, an entity will react differently if there is an event on a public-facing web server versus an event on the energy management system.

For purposes of this report, the joint team considered an "event" as any observable occurrence in a system or network. Events do not necessarily have adverse consequences; they may be innocuous, such as an alert of a user attempting to connect to a file or a web server. These events may not be malicious; however, they could initiate an investigation.

An "incident" is actively violating or purposely intruding on a system or network that could jeopardize the confidentiality, integrity, or availability of the system or network. Incidents may include but are not limited to attempts (either failed or successful) to gain unauthorized access to a system or unwanted disruption or denial of service.[12]

[10] *See* Reliability Standard CIP-005-5 (Cyber Security — Electronic Security Perimeter(s)), https://www.nerc.com/pa/Stand/Reliability%20Standards/CIP-005-5.pdf.

[11] NERC, GridEx, https://www.nerc.com/pa/CI/CIPOutreach/Pages/GridEx.aspx

[12] The joint team's generic definition of incident is somewhat broader than the definitions of Cyber Security Incident and Reportable Cyber Security Incident found in the Glossary of Terms Used in NERC Reliability Standards. NERC, Glossary of Terms Used in NERC Reliability Standards, https://www.nerc.com/files/glossary_of_terms.pdf.

In preparing this report, the joint team also distinguished between the operational environment and corporate environment because of the different risks and impacts of those environments to the overall operation of the participant.

The CIP Reliability Standards definition of “incident” is specific to the reliability of the BES. The joint team observed that participants used the NERC and generic definitions below to fit both environments:

Cyber Security Incident (NERC)[13]

A malicious act or suspicious event that:

- Compromises, or was an attempt to compromise, the Electronic Security Perimeter or Physical Security Perimeter or,
- Disrupts, or was an attempt to disrupt, the operation of a BES Cyber System.

Reportable Cyber Security Incident (NERC)

- A Cyber Security Incident that has compromised or disrupted one or more reliability tasks of a functional entity.

Cybersecurity Incident (generic)

- An unauthorized, unlawful, or unwelcomed action on a computer system and/or network that poses a threat.

IRR Plan

An IRR plan describes how a utility responds to a cyber incident that could operationally and/or financially damage the entity and includes phases and procedures. Establishing clear procedures for handling incidents is a complex undertaking and, though individualized to an organization’s mission, size, structure, and functions, generally contain common elements. IRR plans generally: (1) define their scope (to whom they apply, what do they cover, and under what circumstances); and (2) define computer security events and incidents, staff roles and responsibilities, levels of authority for response (e.g. authority to disconnect equipment), reporting requirements, requirements and guidelines for external communications and information sharing, and procedures to evaluate performance.

[13] On June 20, 2019, the Commission approved revised definitions of Cyber Security Incident and Reportable Cyber Security Incident, which will become effective on January 1, 2021.

III. Observations

The NIST Incident Response Life Cycle identifies the four phases of the incident response process: (1) Preparation, (2) Detection and Analysis, (3) Containment and Eradication, and (4) Post-Incident Activity.[14] The joint team used the NIST Incident Response Life Cycle to organize its observations of participants' responses to the cyber intrusion process. The joint team's observations are discussed below for each of the four phases.

A. Preparation

Preparation involves having incident prevention and response capabilities to ensure that systems and networks are sufficiently secure. These capabilities are typically laid out in IRR plans and include: having appropriate staffing for IRR duties, including clear roles and responsibilities; having adequate procedures and tools necessary for investigation; ensuring team continuity, and professional development; and implementing training that helps personnel recognize malicious activity (e.g., malicious links, odd behavior, or phone calls asking for information).

Observations

The joint team observed that participants approached aspects of preparation differently. With respect to the staffing component of preparation, the joint team found that the importance of well-defined roles and responsibilities became clearer to participants after participating in exercises such as GridEx. The joint team found that the size of the entity, its corporate structure, functions, and geographic footprint influenced the assignment of roles and responsibilities. For example, the roles and responsibilities of smaller entities, because they may only have to respond to an event affecting a handful of systems in one location, are likely to be different from those of large entities who may experience events that affect multiple, geographically dispersed locations involving many systems. A common theme expressed by study participants was the importance of accountability. For example, having a single point of contact to coordinate the event response is beneficial because it reduces the possibility for confusion or miscoordination. Participants also emphasized that it is important for responders to be able to take action based on the severity of the incident without unnecessary delays. For example, some IRR plans authorize operators to disconnect the OT network from the corporate network under certain conditions.

The joint team found that participants are utilizing the Electricity Subsector Coordinating Council's (ESCC) Cyber Mutual Assistance Program.[15] Participants in the Cyber Mutual Assistance Program stand ready to assist in the event an entity asks for help during and

[14] NIST Special Publication 800-61 at 21.

[15] ESCC, Cyber Mutual Assistance Program, https://www.electricitysubsector.org/en/CMA.

following a cyber event. Many participants have contractors on retainer as a precaution to augment their personnel in the event of a major incident, recognizing that responding to some incidents may drain their current personnel despite their best efforts to maintain adequate staffing.

The joint team observed that participants maintained an adequate defense posture by implementing security procedures and tools for monitoring, managing, and mitigating malicious activity with a focus on preventing an event. The joint team found that to ensure preparedness, participant's leveraged technology, such as Application Whitelisting in the CIP environment.[16] The participants also utilized virtualization for quicker recovery.[17]

Participants also recognized the importance of the human element of preparedness. The joint team observed that participants have implemented strong training programs at all levels of staffing. Participants' training included security awareness. Since knowledge and skills in the cyber field need to be constantly updated to remain current with rapidly changing technologies and security threats, participants recognized that IT and security personnel need to develop professionally (e.g., through certifications, software/hardware training).

The joint team observed that most participants conduct quarterly awareness training for all business units. Participants also completed training and cyber exercises at least annually that involve cyber staff from various business units throughout the organization, this is a requirement for the CIP Reliability Standards. The joint team found that it was common for participants to disseminate periodic notices of cyber threats throughout their organizations. Moreover, participants conducted targeted drills such as spear phishing and social engineering exercises. As a result of these efforts, participants reported that their personnel are aware of many types of malicious activity and report suspicious events frequently. Many participants also reported providing incentives (e.g., time off awards, gift cards) to employees that recognize malicious activity during an unannounced exercise.

Apart from drills, the joint team observed that participants constantly develop their personnel's expertise to keep current with technology, vulnerabilities, and threats. Their personnel attend conferences, take classes, participate in working groups, conduct weekly scheduled meetings to discuss emerging threats and vulnerabilities, and develop methods to hire and maintain quality personnel. The joint team found that the emphasis on developing talent from within was in part due to a shortage of cybersecurity and IT professionals. In some cases, participants noted that vacancy announcements remained open for extended periods of time and in some rare cases were not filled. The joint team found that the desire for participants to hire personnel that understand the electric industry and industrial control systems made it even more difficult to find qualified personnel, so entities tend to develop talent from within.

[16] Application whitelisting is a list of applications and application components that are authorized for use in an organization. NIST Special Publication 800-167, Guide to Application Whitelisting (October 2015), https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-167.pdf.

[17] *See generally* NERC, Virtualization and Future Technologies (April 2019), https://www.nerc.com/pa/Stand/Project 201602 Modifications to CIP Standards RF/Project 2016-02_Virtualization_and_Future_Technologies_Case_for_Change_White_Paper_04182019.pdf.

Moreover, the shortage of qualified personnel is compounded if staff training requests cannot be accommodated due to operational time constraints or lack of staff coverage. The joint team found instances where personnel were registered to take IT or security training but could not due to such constraints.

The joint team observed that participants also prepared by looking closely at the global cyber-threat landscape. Participants performed research and threat assessments across the globe since Internet-connected systems can be reached from anywhere on the planet. For example, all participants had reviewed the recent attacks against Ukrainian systems, the attack methodologies, and discussed how such an attack would have affected their own systems.[18] The joint team observed that when an assessment of the threat and/or vulnerability is conducted, participants considered protecting the OT network a priority. Responding to NERC Alerts is another form of preparation observed by the joint team. An example of such an alert is the October 5, 2017 Industry Recommendation on "Supply Chain Risk," where all participants acted on the NERC Alert.[19] The joint staff learned that when participants receive NERC Alerts, typically the personnel assigned to assess the alert meet to discuss and plan a strategy to address it. The joint team noted the seriousness and diligence taken when NERC Alerts are issued.

Key Take-Aways

- Effective IRR plans contain well-defined personnel roles, promote accountability, and, where appropriate, empower personnel to take action without unnecessary delays.
- Effective IRR plans leverage technology and automated tools while also recognizing the importance of human performance.
- Effective implementation of IRR plans requires well-trained personnel who are constantly updating their skills.
- Effective IRR plans incorporate lessons learned from past cyber events and simulation of real-world events identifying clear shortfalls in the IRR plan.

B. Detection and Analysis

The detection and analysis of a suspicious event is a complex task requiring trained personnel and tools to defend the targeted cyber systems. Moreover, to recognize that an event qualifies

[18] Department of Homeland Security, Alert (IR-ALERT-H-16-056-01) Cyber-Attack Against Ukrainian Critical Infrastructure (February 25, 2016), https://ics-cert.us-cert.gov/alerts/IR-ALERT-H-16-056-01.

[19] NERC, Industry Recommendation: Supply Chain Risks, https://www.nerc.com/pa/rrm/bpsa/Pages/Alerts.aspx.

as an incident, personnel must navigate often ambiguous, contradictory, and incomplete information from the initial event.

Detection includes using security tools to identify network scans. Network scans against the perimeter of an OT network may raise concerns because the network is typically located within the protected environment of an entity's corporate network. A network scan at this location means that the attacker may have penetrated multiple trust-zones before reaching the OT network or related critical assets. IRR plans typically include an escalation process for events depending on the risk to the operational environment. These tools may also be used to detect and analyze more complex events such as known attacks, characteristics of a suspected uncategorized attack, or traffic that indicates a compromise has already occurred within the system.

Event analysis requires a determination of scope, such as which networks, systems, or applications are affected; who or what caused the incident; and how the incident is occurring (e.g., what tools or attack methods are being used, what vulnerabilities are being exploited). Typically, when an entity suspects that an incident has occurred, it rapidly deploys a team to perform an initial analysis to determine the scope of the incident. The initial analysis provides information to prioritize subsequent activities, such as containment of the incident and deeper analysis of the effects of the incident. This can include baselining events, which is the practice of monitoring resources to determine typical patterns so that significant deviations can be detected. Baselining events contributes to the efficient use of personnel because many events are network scans or reconnaissance and thus may not warrant the same response as cyberattacks designed to disrupt the protected networks.

Observations

As discussed in the previous section, participants employ exercises and drills to test their ability to identify and respond to cyber events and incidents. The joint team observed that, to detect evolving and persistent attacks, participants used security tools to perform network scans and test unauthorized services. Participants used tools to scan their systems to determine the security posture of their system and quickly identify vulnerabilities such as software flaws, missing patches, malware and misconfigurations. The security tools used by participants worked across a variety of operating platforms, and they include industrial control systems. Beyond looking for vulnerabilities and missed patches that may increase the risk of compromise, the joint team observed that security tools also assist in baselining systems, as well as in the configuration management process (i.e., by providing an automated reporting process, for example, for information related to ports and services, patching, connected systems, installed software, and baselines). In many cases these automated security tools are used for documentation requirements that would be difficult to perform manually. However, participants recognized that they cannot rely on automated tools only; instead, participants used these tools to complement their personnel's training and expertise. Participants explained that they regularly communicate with IT personnel, management, and third parties to notify, identify, and react to vulnerabilities. They also train and mentor less experienced personnel and remain current on threats and vulnerabilities.

The joint team observed that participants did not have a common definition for the terms “suspicious,” “event,” and “incident.” As a result, participants reacted differently to similar events. For example, some participants made purely subjective determinations that something was suspicious, rather than apply objective criteria based on system configuration, history of network activity and risk. All participants agreed, however, that suspicious events in the corporate environment are more prevalent than in the more critical OT network. The joint team noted that many participants established risk thresholds for performing analysis. While participants’ risk thresholds varied, a common technique employed the use of flow-charts or decision trees to determine quickly when a predefined risk threshold is reached, and a suspicious set of circumstances qualifies as an event. These techniques can be integrated into the configuration of the security tools used across the system.

The joint team found that a suspicious event will elicit a different reaction depending on the affected trust-zone. For example, personnel will react quickly when they detect suspicious events in the higher trust zone of the operational environment (i.e., within an Electronic Security Perimeter) to ensure there is no malicious activity. By contrast, the joint team observed that personnel may respond differently to a suspicious event in a trust-zone that is directly connected to the Internet because, by design, the Internet-facing trust-zone is segregated from the operational environment.[20] Both suspicious events would be investigated by personnel, but follow up on the OT network event would most likely have more urgency because of the higher risk. The joint team expected this response because those actions are based on risk and impact to system operations.

The joint team observed that the participants align the use of automated security tools for detection with their event and IRR plans to determine risk thresholds and create alerts. Participants configured tools based on industry best practices, cybersecurity standards, recommended settings by the vendor, and the experience of personnel maintaining the system. The joint team discussed the various tools that assist with detecting and analyzing known cyber events and patterns automatically, and participants indicated that the tools produced results that gave security personnel staff a high degree of confidence.

As emphasized above, personnel training is critical to effective detection and analysis. The joint team observed that personnel had a thorough understanding of the operating environment, were able to recognize suspicious events, and could distinguish between a benign and suspicious event.

[20] For example, a firewall at an Internet gateway will often be probed constantly for weaknesses because the system is directly connected to the Internet. In this setting, detecting something new would likely qualify as suspicious, and the security team would investigate to ensure it is benign or low risk. A new type of suspicious event is difficult to discover because of the volume of unauthorized attempts directed at systems connected to the Internet. By contrast, a firewall located at the perimeter of the operational network environment (e.g., an electronic security perimeter) should detect very little unauthorized activity within this trust-zone. If an unauthorized or unexpected scan is conducted against this firewall, the security team would react differently from a firewall located at the Internet gateway because of the risk and profile of the expected traffic at each location (i.e., the systems at the electronic security perimeter should not detect suspicious traffic).

Key Take-Aways

- Baselining is an effective resource utilization tool that allows personnel to detect deviations from normal operations.
- Flow-charts or decision trees are useful tools to determine quickly when a predefined risk threshold is reached, and a suspicious set of circumstances qualifies as an event.

C. Containment and Eradication

Containment is the process of ensuring that an incident's scope or impact is limited. Once a suspicious event is verified as an incident, entities implement a plan to contain the malicious activity. To implement containment, the entity's risk assessment should reflect an understanding of the impact of potentially losing a system or device in the operational environment. For example, entities typically have documented processes that address who has the authority to decide whether to shut down an operational network or allow the network to remain functional (e.g., to monitor the attacker before taking action).

Containment is not always the best option. For example, if the affected system is disconnected to prevent malware from communicating with the attacker, this can result in predefined destructive actions by the malware, such as erasing data or rendering equipment unusable. Another important part of containment is effective communication amongst security personnel, corporate leadership, and relevant third parties to ensure the impact to the operational system is well understood regardless of the containment strategy.

Eradication involves removing the threat from the impacted systems in an attempt to return the system to a normal, functioning state. The eradication process is often challenging because an incident response team cannot guarantee that it has identified all the systems that were compromised. For example, if the attacker has been in the system for an extended period of time, the attacker potentially compromised many systems and gained a foothold across systems in the corporate and operational environments. The attacker may also create false credentials and/or installed backdoors to allow for future access. As with containment, communication is an important part of eradication. Ineffective communication provides the attacker with an advantage because the attacker can exploit any confusion during the IRR process. Because a widespread compromise has potential consequences for the entire organization, good communication across the entity is critical to manage such an event and defeat the attacker.

Observations

The joint team observed that participants utilized various containment strategies in their IRR plans. Strategies ranged from monitoring malicious traffic to better understand it, to unplugging and isolating the operational network to protect reliable operations. Participants indicated that their response depended on the nature of the event, or incident, and the expected impact to the BES. Accordingly, IRR plans typically included multiple options to respond to an incident based on the severity of its expected impact on reliability.

Containment strategies should balance the need for responsiveness with the potential of operational impact if a containment strategy, such as isolation, is performed too soon. For example, in discussions with participants, personnel stated that they make assessments quickly and take actions based on initial assessments, but these actions may later be determined to have not been the best approach.   The joint team found that some participants authorized personnel to take independent action in response to an event, such as disconnecting communications between the operational and corporate networks, to contain incidents with the potential for serious impacts. As an example, an operator may have the authority to assess the situation and sever service to non-essential networks, such as the corporate IT network, leaving the OT network functioning, in effect islanding the operational networks.[21] Participants explained that conducting tests and training is key to this approach so that security teams know how OT networks react when such action is taken.

The joint team learned through discussions with participants that disconnection is considered containment, not eradication, because it isolates the operational network and prevents further external communication to the attacker. Some participants indicated that a decision to disconnect may be based on malicious activity outside of the operational network, thereby protecting critical operational systems before malicious activity is detected within the operational network. This illustrates the importance of broad situational awareness across an entity to maintaining operations and protecting critical systems in the operational network. This also demonstrated to the joint team the importance of good communication between security teams, business units, and leadership.

Once the system is isolated, the security team may be able to take actions to ensure that evidence of the event is preserved and to begin the eradication process. Participants explained that detecting malicious activity does not always mean a single compromise in the system. Rather, the detection could be an indicator of a larger compromise. This reinforces the importance of evidence collection and continued analysis. However, participants indicated that in some cases safety, health, and operational impact may take precedence over the need to preserve evidence.

The joint team observed that some participants use technology to react more rapidly. For example, many participants use virtualization, which is the use of software to operate as if it were actual physical hardware, in the corporate and the operational environments. By virtualizing hardware, one physical device houses many virtual devices, thus reducing costs in hardware and building space. Since a virtualized device is software that can be easily saved and restored, having virtualized systems provides a method to take a snapshot of important forensic information such as memory contents, connection information, date and operating state. Once this important forensic information is saved, a safe copy of the virtualized system can be restored in minutes.   One participant explained that having a virtualized device in its operational environment saved hours of work when a software glitch occurred requiring multiple machines to be reloaded. The participant stated that the same response would be

[21] Participants considered disconnection containment, not eradication, because it isolates the operational network and prevents further external communication to the attacker.

used if they had to reinstall a new machine because of a cyberattack and that a virtualized device would potentially save hours of installation time.

The joint team and participants discussed scenarios where a compromised device or system is the source of the attack and must be contained to prevent further compromise, damage, or exfiltration of data. Once the compromised devices or systems are contained, which may require their removal from the network, participants stated that they begin the preservation of evidence and recovery processes. To do this, participants indicated that they may assemble special teams, depending on the level of the perceived risk to the system; and they emphasized the importance of communication among personnel during the entire incident.

In responding to various cyber attack scenarios, important factors identified by participants include the length of time the attacker has been in the system, as well as the type of attack. For example, if the attacker has established a “command-and-control” channel, it may be better for the participant to understand what the attacker is doing before taking any containment and eradication actions. Based on the security team’s analysis of the attack, the team would develop a plan on how best to contain and eradicate the attacker. This is a common strategy for a system suspected to have been compromised for a long period of time or if the security team suspects that severing the connection may trigger wiper malware (e.g., malware to destroy the affected system by wiping all files and operating system). The joint team observed that IRR plans do not necessarily specify how to approach this type of event; rather the IRR plan may only set out the conditions when to monitor and when immediate containment is necessary.   In any case, the IRR plan should consider the possibility of destructive malware being activated if the attacker’s connection is severed.

The joint team found that participants preserve and analyze evidence beyond the compromised systems. Gathering network evidence is an important part of that collection process because it aids in determining the scope of the compromise, as well as its duration. Each participant discussed its approach, the chain of custody, the level of evidence preserved, and how its IRR plan addressed the scope of evidence preservation. During the evidence gathering process, the joint team learned that it is not uncommon to discover additional malware once the evidence is examined closely. Some participants cautioned that malware determined to be unrelated to the incident and that poses little risk may distract from the true mission of response and recovery by diverting valuable time and resources.

Several participants discussed the legal implications of evidence collection and retention. Participants stated that their legal departments participate in developing their IRR plans because applicable systems may contain personal data or company sensitive data requiring special protection. Additionally, an entity must plan for the chain of custody for evidence and how the evidence could be used for future training and exercises. Participants coordinated with their legal departments so that the security teams understood their legal obligations; for example, in the recovery phase it is important to know how long to preserve the evidence for legal and compliance requirements.

Knowing how long eradication will take is inherently difficult because an incident may be more extensive than when initially discovered. Accordingly, participants discussed how they budgeted and planned resources for responses of potentially indeterminate length. IRR plans account for this by having enough detail to allow for cost estimates and resource needs for events/incidents of short duration; but IRR plans are flexible enough that they do not restrict the team from effectively addressing incidents that go on for months or years, which are difficult to budget for. Participants raised scenarios where an incident lasts for an indeterminate time and how such an event would stress an entity's resources. Prolonged incidents may strain personnel and affect employee retention, which is a serious concern given that continuity of personnel and knowledge is vital to reliable operations. To address incidents with long time horizons, the IRR plan may call for creation of a dedicated response team with subject matter experts, legal staff, and contractors trained in IRR. Regardless of the severity of an incident, participants explained that establishing a team lead for the incident response is vital. Participants also stated that this type of scenario is usually costly and other projects will most likely be put on hold while a large incident is fully addressed.

Key Take-Aways

- If an IRR plan containment strategy includes islanding operational networks, there should be a thorough understanding of the potential impact of such a decision (e.g., through tests and training).
- IRR plans should consider the possibility that a containment strategy may trigger predefined destructive actions by the malware.
- Evidence collection and continued analysis are important to determine whether an event is an indicator of a larger compromise.
- IRR plans should consider the resource implications of incident responses of indeterminate length.

D. Post-Incident Activity

Post-incident activity encompasses recovery from an incident and documenting and analyzing the incident to prepare lessons-learned reports. Recovery involves restoring the system to an operational state.

Lessons-learned reports address all aspects of an incident response, including the response techniques used, system shortfalls, system strengths, training and exercises, future event response, budgeting, and staffing. Lessons-learned reports are reviewed and used by the teams that responded to the incident, as well as personnel in the entity responsible for reliability of operations, leadership, and budgets. Lessons-learned reports not only discuss what went wrong, but also what was done correctly. Not all incidents warrant a lessons-learned report, however. A lessons-learned report is a substantial task, and typically there is a threshold for the types of events that warrant a report. For example, a virus detected and quarantined by anti-virus software may be classified as an event, with its details preserved for future reference, but it likely does not warrant a formal lessons-learned report.

Observations

The joint team observed that it is possible for there to be tension between incident response and documenting incidents to develop lessons-learned reports. For example, documenting an incident that is ongoing could divert resources from the incident response process and delay restoration of systems to an operating state. However, participants agreed that developing lessons-learned reports was a necessary post-incident activity.   Moreover, large-scale exercises with high-consequence incidents reinforced to participants the need to document their actions to ensure that nothing is missed, as well as to provide details for lessons-learned reports. One practice identified by participants to use resources efficiently is designating focused resources to document the response process. Additionally, participants used tools to assist with the process, such as ticketing software and incident response management tools.[22]

Key Take-Away

- Effective IRR plans implement lessons-learned from previous incidents or training exercises.

IV. Other Topics

In addition to the observations based on the NIST life cycle, the joint team provides the following additional observations on the participants' IRR plans, testing of the IRR plans, how participants report incidents, and verbal communications.

A. IRR Plans

The participants organized their IRR plans differently. Some participants have an IRR plan that covers the operational network and a separate IRR plan that covers the corporate network. Other participants have a single plan (or mirrored plan) that spans the operational and corporate networks. The joint team found that large companies tend to have two or more IRR plans that address different network environments such as operational and corporate networks. The IRR plans used a standard framework that was modified to address the specific business units' needs.

For entities with more than one IRR plan, the joint team found that the operational network IRR plans often follow the reporting requirements in the CIP Reliability Standards. The CIP Reliability Standards address cyber events (i.e., Cyber Security Incident), but they are reported differently from cyber incidents discovered in the corporate network. Some participants use the NERC definition of Reportable Cyber Security Incident as a trigger to initiate IRR plans, and they view any cybersecurity condition as a "cyber event" if the definition does not meet the "reportable" threshold. This approach tends to reduce the need to revise IRR plans because incidents are not as frequent.

[22] Ticketing software is a method used to track issues within an organization; for example, an IT helpdesk may assign ticket numbers to issues that are called in to them as a means to track and ensure resolution.

The joint team observed that when participants use one IRR plan, the plans are reviewed more frequently. These IRR plans are divided into two sections: event and incident. The event section of the IRR plan is updated more frequently because everything starts as an event and as a result there are a larger number, in most cases an event results in system or human error and is resolved. The event will escalate to an incident when malicious activity is identified resulting in a smaller subset of events. Moreover, entities with one IRR plan use it more often since it is intended for the entire organization and involves multiple networks with various trust levels.

Some participants used one incident response team that covered all business units. One participant explained that it previously had two incident response teams: one for the corporate network and one for the operational network. The participant explained that the entity combined them into one team after a hardware failure lead to communication, authority issues, and lack of system awareness across the business units. Following a lessons-learned report, the participant determined that it could have responded more efficiently if a single team worked across the entire organization. The participant stated that following the consolidation, training exercises went more smoothly, and personnel had a better understanding of the operating conditions of the entire entity.

To test IRR plans, participants use the requirements of the CIP Reliability Standards, which have a 15-month testing interval requirement. Many participants test corporate and operational systems on a 12-month cycle. Additionally, entities conducted smaller exercises more frequently to ensure their process, procedures, and IRR plans met emerging threats. Many entities engage in large group exercises, such as GridEx, to test and exercise their IRR plans. For entities that used GridEx to exercise their plan, all participants stated the exercise was helpful and many stated that they modified their IRR plans after the exercise was completed.

All the participants have a formal process for updating their IRR plans. Typically, the participants have regular cybersecurity meetings that include the review of these plans. For example, when a new threat or vulnerability is discovered, participants review their IRR plan to determine if it adequately addresses the new threat or vulnerability. The joint team observed, however, that it is uncommon to change IRR plans simply because a new threat or vulnerability is discovered unless the risk is high and thought to be persistent.[23]

Many of the participants' IRR plans use flow-charts to make a quick determination of when an event becomes an incident. Flow-charts are a useful tool for quickly aiding in the decision process. A basic flow-chart may require a person with more experience to use professional judgment in decision-making process. Conversely, if the flow-chart is too complex, the process may become overly complicated and difficult to implement. Participants employing flow-charts

[23] An example of a vulnerability that may generate a change to an IRR plan is the "SPECTRE and MELTDOWN" chip vulnerability. Department of Homeland Security, Alert: Meltdown and Spectre Side-Channel Vulnerability Guidance (Jan. 4, 2018), https://www.us-cert.gov/ncas/alerts/TA18-004A.

indicated that they had used them for a few years and that they were updated based on testing/events to allow for better process flow of an event and escalation to an incident.

Some participant's IRR plans include using a third party with experience responding to complex compromises involving multiple devices and systems. Participants that use third parties indicated that it was important for the outside personnel to understand in advance the participant's system and operational needs prior to an event. Additionally, the third-party incident response team can help develop training exercises and assist with the development or maintenance of IRR plans.

B. Reporting Incidents

Reporting incidents accurately and quickly is critical to an entity's response efforts. Additionally, reported information may provide other entities with valuable information so they can prepare their system to detect an attack. Reporting also benefits those that might already be compromised and are not aware of the event by providing them a starting point to look for the malicious activity. There are also governmental agencies, such as the Department of Homeland Security National Cybersecurity & Communications Integration Center (DHS-NCCIC), that tracks events and incidents, as well as private organizations that assist with reporting and incident response, such as the Electricity Information Sharing and Analysis Center (E-ISAC).

Participants all agreed that sharing information is crucial. Participants indicated that they report suspicious events and incidents in their corporate and/or operational networks to the E-ISAC, DHS-NCCIC, and/or other security organizations to inform the wider industry of malicious activities.[24] Participants also share information regarding suspicious activity directly with one another in an informal manner, because getting ahead of a potential incident is not only helpful for their organization, but also for the utility sector as a whole. Participants commented that they would like to send and receive information in an automated way because it takes time to review all the information when they receive it manually.

Participants considered social media activity as a potential source of suspicious activity. The joint team observed that participants look at social media to see if information was posted that could be used to exploit their system. For example, an employee might post a whitepaper on the energy management system used by its company. Under this scenario, the entity may examine the details of the whitepaper and make a determination if the activity warrants security enhancements such as monitoring for conditions that may have been discussed in the posted whitepaper.

Participants also consider phone calls, mail, and e-mail as potentially suspicious. One example is a spear phishing email that asks for specific information about operating conditions

[24] Under Reliability Standard CIP-008-6 (Cyber Security — Incident Reporting and Response Planning), which becomes effective on January 1, 2021, responsible entities must notify the E-ISAC and, if subject to the jurisdiction of the United States, DHS-NCCIC, or their successors, of a Reportable Cyber Security Incident and a Cyber Security Incident that was an attempt to compromise.

and/or systems. Participants addressed these types of activities in their cyber training to employees. Awareness of this type of attack means that personnel are less likely to fall prey to them in both the corporate and the operational network environments.

C. Communication

Participants described the importance of relationships and establishing effective communication among an organization's departments (e.g., human resources, legal, management), as well as with outside groups (e.g., partner incident response teams, law enforcement). Policies and procedures define the structure of cybersecurity programs; but without good communication, entities may not implement them effectively, which may lead to lapses or wasteful duplication of effort.

The joint team observed that participants recognized this by establishing in their IRR plans a clear communication path with the management personnel authorized to make decisions in response to an event. For example, the joint team found that the participants' communication paths remained open for those on call 24 hours a day, seven days a week even if the security operations center was not open during that time.

Developing one IRR plan for an organization may be more conducive to effective communication because, under a single IRR plan, one team monitors corporate and operational network cybersecurity. Accordingly, one team will have situational awareness of cybersecurity concerns across all network environments. This model, however, may be difficult for large entities because of the number of cyber assets and personnel that the IRR plan must cover. In addition, for entities that are divided into multiple companies, it may not be feasible to have one team and/or IRR plan. Another potential drawback is that the security may not have enough expertise in the control system environment simply because the network and number of assets the team is responsible for is so large. One approach discussed by participants is to have specialized sub-teams with specific expertise and clear communication protocols. These specialized security teams may be, for example, a firewall, network, energy management, or Active Directory team.

Maintaining secure communications between network environments is another important consideration. Participants discussed communication using a common Security Information and Event Management (SIEM)[25] system for the corporate and operational networks. Specifically, logged information is passed securely from the CIP environment to the corporate environment where the SIEM is located and all information is correlated at a central location. Using a central location allows for full view of all the cyber environments. When this approach is taken, personnel monitoring alerts are trained to understand the significance of the alert coming from the CIP environment. The joint team observed that there is no one superior SIEM

[25] Security Information and Event Management (SIEM) is technology that supports threat detection and security incident response through the real-time collection and historical analysis of security events from a wide variety of event and contextual data sources. It also supports compliance reporting and incident investigation through analysis of historical data from these sources. The core capabilities of SIEM technology are a broad scope of event collection and the ability to correlate and analyze events across disparate sources.

design; the optimal design depends on the entity's networked environment design. What is important is communication between these business units and that leadership is aware of the conditions across all of their system.

Key Take-Away

- Using a single or similar IRR Plan between business units may provide a better understanding of the IRR process across the entity.

V. Conclusion

IRR plans are important resources for addressing cyber threats, and effective IRR plans can mitigate the natural advantages that cyber attackers possess. Because attackers operate covertly to gain footholds across networks, effective IRR plans should be in place and response teams should be prepared to detect, contain, and, when appropriate, eradicate the cyber threat before it can impact the utility's operations. Recognizing that there are differences among utilities, there is no one best IRR plan model. However, based on the joint team's observations, there are practices that utilities should consider when developing their own IRR plans.

Appendix 1 – Glossary of Terms Used in Report

Application whitelisting: A list of applications and application components (e.g., libraries, configuration files) that are authorized to be present or active on a host according to a well-defined baseline. Application whitelists are enforced by application whitelisting programs, also known as application control programs and application whitelisting technologies.

Baselining: Monitoring resources to determine typical utilization patterns so that significant deviations can be detected.

BES Cyber Asset (BCA): A Cyber Asset that if rendered unavailable, degraded, or misused would, within 15 minutes of its required operation, misoperation, or non-operation, adversely impact one or more Facilities, systems, or equipment, which, if destroyed, degraded, or otherwise rendered unavailable when needed, would affect the reliable operation of the Bulk Electric System. Redundancy of affected Facilities, systems, and equipment shall not be considered when determining adverse impact. Each BES Cyber Asset is included in one or more BES Cyber Systems. A Cyber Asset is not a BES Cyber Asset if, for 30 consecutive calendar days or less, it is directly connected to a network within an ESP, a Cyber Asset within an ESP, or to a BES Cyber Asset, and it is used for data transfer, vulnerability assessment, maintenance, or troubleshooting purposes.

Command-and-control channel (C2C): The communications channel that is established between the compromised computer and the command-and-control server.

Cyber Alert: Notification that an event (or multiple similar events) have occurred.

Cyber Assets: Programmable electronic devices, including the hardware, software, and data in those devices.

Cyber Event: An observable occurrence or change in a cyber system requiring investigation. A cyber event may escalate into a cyber incident but not always.

Cyber Mutual Assistance Program (CMA): Program that is an industry framework developed at the direction of the ESCC to provide emergency cyber assistance within the electric power and natural gas industries. The CMA Program is composed of industry cyber experts who can provide voluntary assistance to other participating entities in advance of, or in the event of, a disruption of electric or natural gas service, systems, and/or IT infrastructure due to a cyber emergency. As the CMA Program develops, additional initiatives will be considered and implemented based on the needs and input of the entities participating in the CMA Program.

Electronic Security Perimeter (ESP): The logical border surrounding a network to which BES Cyber Systems are connected using a routable protocol.

Gateway: A network device that is a central ‘chokepoint’ for inbound and outbound traffic from one network to another different network. It is possible to have multiple gateways.

GridEx: GridEx is an electric sector side exercise for utilities to demonstrate how they would respond to and recover from simulated coordinated cyber and physical security threats and incidents, strengthen their crisis communications relationships, and provide input for lessons learned. The exercise is performed every two years.

Inbound traffic: Network traffic that is coming from one trust-zone and moving into a network with a different trust-zone and usually passing through a gateway.

Indicator of compromise (IOC): A forensic remnant of an intrusion that can be identified on a computer or network device that with high confidence an intrusion has occurred by examining the forensic data. An IOC signature is created to proactively monitor a system for malicious activity.

Internet facing: A device that is directly connected to the Internet. These devices typically can be identified by anyone that is connected to the Internet, and as a result are susceptible.

Islanding: Disconnecting computer networks so that there is no connection between the two networks (e.g., there is no possibility for inbound or outbound traffic from the islanded network). The only method to move information into an isolated network is to use physical media such as USB drive or some type of removable media. Also referred to “air gapping.”

Malware: Malicious code or software that is designed to perform an unauthorized action on a computer system.

Outbound traffic: Network traffic that is leaving one trust-zone to another level trust-zone and usually passing through a gateway.

Physical machine: A term that is used to differentiate a physical machine (e.g., computer) from a virtual machine. The physical machine is tangible hardware that occupies space.

Physical Security Perimeter (PSP): The physical border surrounding locations in which BES Cyber Assets, BES Cyber Systems, or Electronic Access Control or Monitoring Systems reside, and for which access is controlled.

Security Information and Event Management (SIEM): Technology that supports threat detection and security incident response through the real-time collection and historical analysis of security events from a wide variety of event and contextual data sources. It also supports compliance reporting and incident investigation through analysis of historical data from these sources. The core capabilities of SIEM technology are a broad scope of event collection and the ability to correlate and analyze events across disparate sources.

Security Operations Center (SOC): An operations center focused on cyber security where a Security Information and Event Management system would most likely be found.

Spear phishing: A malicious email that is sent to a recipient that appears to come from a trusted sender. Typically, these emails are sent with the intent to perform a malicious action on the recipient’s computer.

Suspicious event: A confirmed cyber event that compromises or attempts to compromise a cyber system and disrupts or attempted to disrupt the cyber system.

Trust-zone: A network segment where the data that is processed and traverses within the network segment is trusted.

Virtualization: The concept of converting physical hardware into software. The software operates as if it was the actual physical hardware. Virtualization allows one physical device to encapsulate many virtual devices and thus reduces costs in hardware and building space.

Virtual Machine: Hardware that has been virtualized into software. Many virtual machines can be installed on a single physical machine.

Wiper malware: Destructive malware that erases the operating components of a cyber device rendering the cyber device unusable.

Appendix 2 – Acronyms Used in Report

BES	Bulk Electric System
CIP	Critical Infrastructure Protection
CYPRES	Cyber Planning for Response and Recovery Study
DHS-NCCIC	Department of Homeland Security National Cybersecurity & Communications Integration Center
E-ISAC	Electricity Information Sharing and Analysis Center
ESCC	Electricity Subsector Coordinating Council
FERC	Federal Energy Regulatory Commission
GridEx	NERC Grid Security Exercise
IRR	Incident Response and Recovery
IT	Information Technology
LAN/WAN	Local Area Network/Wide Area Network
NERC	North American Electric Reliability Corporation
NIST	National Institute of Standards and Technology
OT	Operational Technology
REs	Regional Entities
SIEM	Security Information and Event Management

2020 FERC, NERC and REs Report

Cyber Planning for Response and Recovery Study (CYPRES)

Federal Energy Regulatory Commission
North American Electric Reliabity Corporation
Regional Entities
September 2020