CCE Phase 1: Consequence Prioritization

Consequence-driven Cyber-informed Engineering

Idaho National Laboratory

CYBERCORE integration center

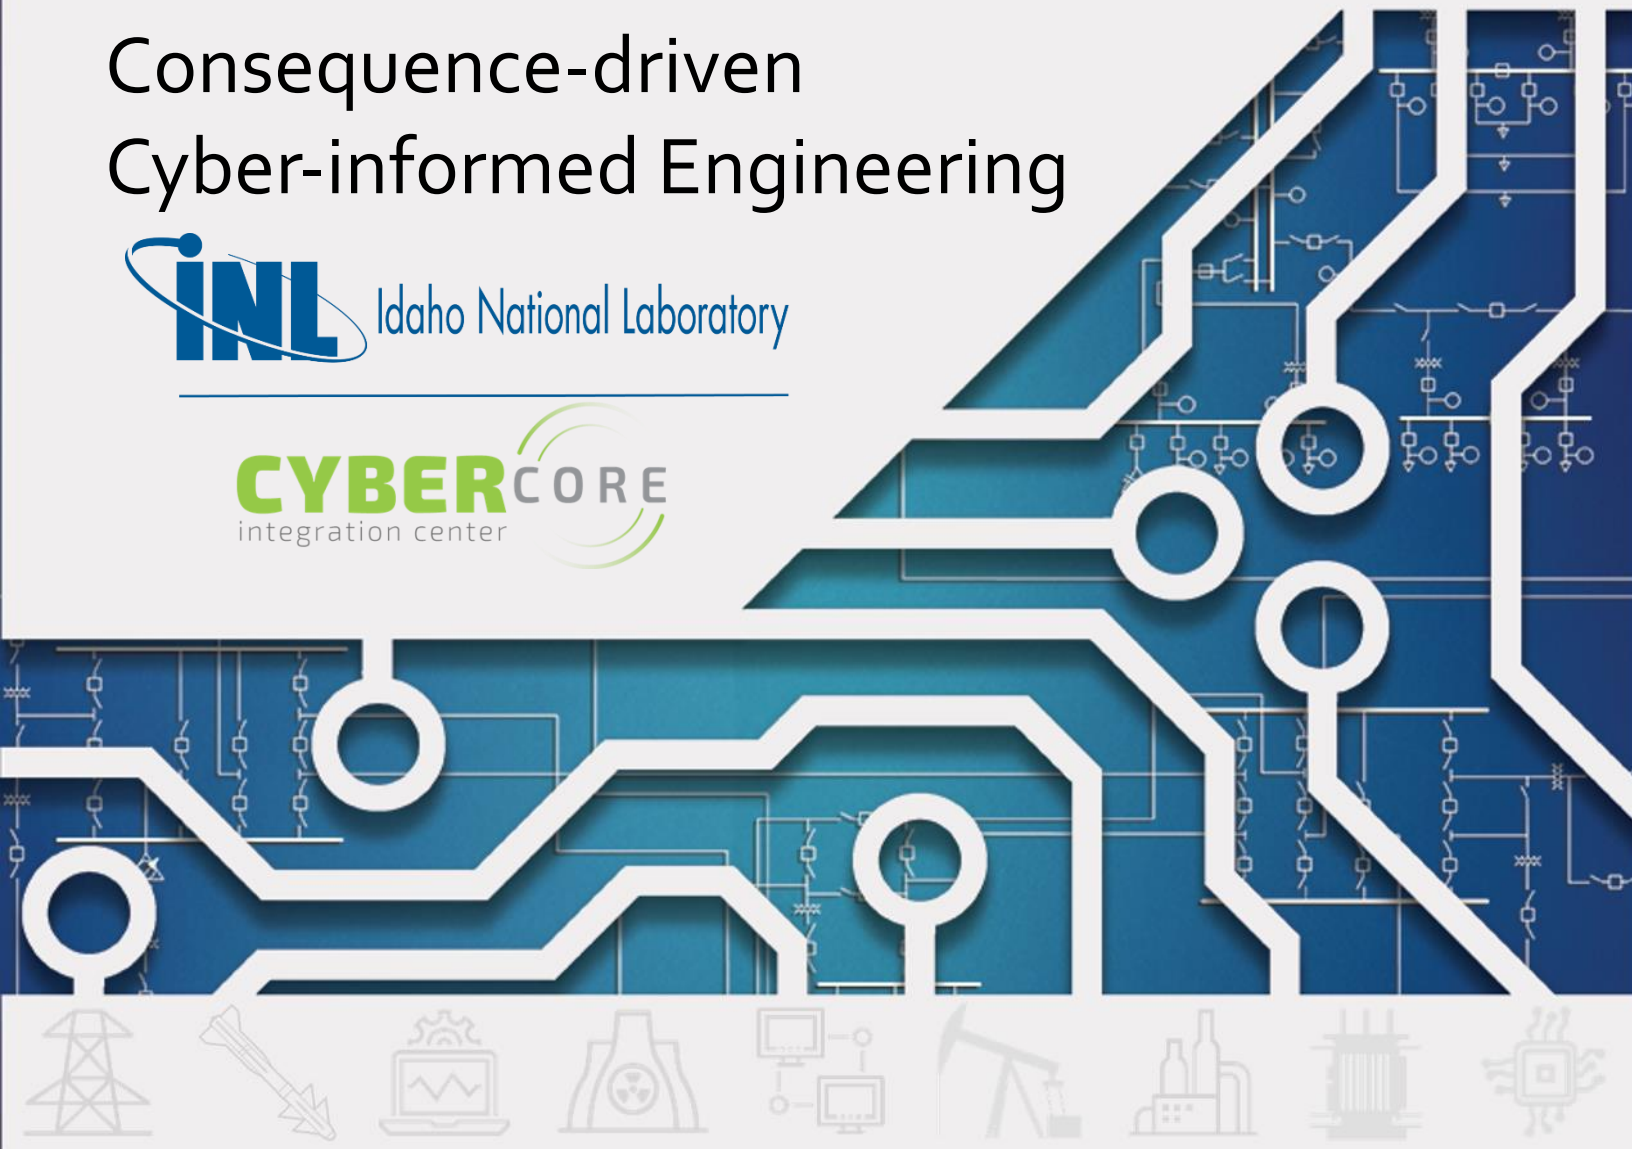

Prepared By: Sarah G. Freeman, Nathan Hill Johnson, and Curtis P. St. Michel
Cybercore Integration Center
Idaho National Laboratory

May 5, 2020

INL-EXT-20-58089

CCE Phase 1: Consequence Prioritization

Consequence-driven Cyber-informed Engineering

Sarah G. Freeman
Control Systems Cybersecurity Analyst

Nathan Hill Johnson
Control Systems Cybersecurity Analyst

Curtis P. St. Michel
Cybercore Technical Director

Idaho National Laboratory
Cybercore Integration Center
Idaho Falls, Idaho 83415

http://www.inl.gov

Prepared for the
U.S. Department of Energy
Office of National & Homeland Security
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517

Page intentionally left blank

CCE Phase 1: Consequence Prioritization

Introduction

Idaho National Laboratory (INL) developed the Consequence-driven Cyber-informed Engineering (CCE) methodology to provide public and private organizations with steps to work collaboratively and establish a working relationship to protect critical infrastructure and other national assets. This process is a considerable undertaking, iterative in nature, and—as time and resources allow—should become a part of a company's culture. By focusing on the impact of potentially negative Events, CCE provides a better understanding of how and why adversaries can affect critical functions and services using cyber-enabled sabotage.

The CCE methodology consists of four phases:

Phase 1: Consequence Prioritization

> During this phase, the CCE Team works together to develop the boundaries and thresholds for Events and cyber-Events that could be catastrophic to the organization. They are then prioritized to determine which can be deemed High Consequence Events (HCEs).

Phase 2: System-of-Systems Analysis

> Here the team maps out the systems and processes related to the HCEs identified in Phase 1, and then investigates the dependencies and "unverified trust" which would enable them.

Phase 3: Consequence-based Targeting

> The team refines and develops the targeting requirements an adversary would need to fully understand the attack in detail and, consequently, carry it out.

Phase 4: Mitigations and Protections

> In the final phase, the priority is to take the possibility of the physical effect through cyber means out of the equation using engineering or process changes. If this is not possible, use the detailed targeting requirements developed during Phase 3 to detect adversary activity and implement other types of mitigations.

Consequence Prioritization

This document describes the process for Consequence Prioritization, the first phase of the CCE methodology. The primary goal of Consequence Prioritization is to identify potential disruptive cyber-Events—that is, physical Events that are achievable through cyber means—that would significantly inhibit an organization's ability to provide the critical services and functions deemed fundamental to their business operations or mission.

These disruptive cyber-Events, defined as High Consequence Events (HCE), could include failures or natural disasters, but they should also include cyber misuse of systems and the unique digital dependencies of critical infrastructure assets. While other efforts have been initiated to identify and

mitigate disruptive cyber-incidents at the national level, such as Presidential Policy Directive 41,[a] this process is intended to be used by individual organizations to complement those efforts.

Described another way, Consequence Prioritization considers threats greater than those addressable by standard cyber-hygiene and includes the consideration of events that go beyond a traditional continuity of operations (COOP) perspective.

Finally, Consequence Prioritization is most successful when organizations adopt a multi-disciplinary approach, engaging both cybersecurity and engineering expertise, as in-depth engineering perspectives are required to recognize, characterize, and mitigate HCEs. Figure 1 provides a high-level overview of the prioritization process.

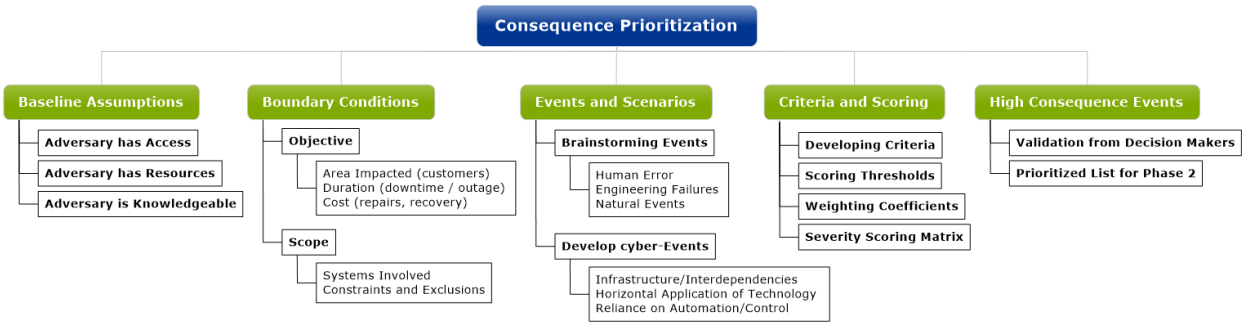

Figure 1: CCE Prioritization method overview.

Establish Baseline Assumptions

Baseline Assumptions:

- **Access has been achieved**
  - Adversary has logical and physical access, including all credentials, IP addresses, firewall and application access, distribution management system (DMS) access, distributed control system (DCS) access, etc.

- **Adversary is knowledgeable**
  - They understand critical equipment and processes and possess the knowledge required to impact the system.

- **Adversary is well-resourced**
  - They have access to the required equipment, engineering expertise, and tools.

[a] President Barack Obama's Presidential Policy Directive 41, "United States Cyber Incident Coordination," July 26, 2016, can be found at https://obamawhitehouse.archives.gov/the-press-office/2016/07/26/presidential-policy-directive-united-states-cyber-incident.

Objective, Scope, and Boundary Conditions

The CCE Team's first step should be to formally establish and finalize the Objective, Scope, and Boundary Conditions for the CCE engagement. These scoping tasks help better define the area or scale of interest. These concepts must revolve around the critical functions and services that the organization provides. These critical functions and services make up the purpose or mission of the company or organization, and they often have a direct impact on the community or nation. For a large organization which provides services deemed essential to national interests, those interests often become part of the Boundary Conditions.

Rather than focus on some aspects of likelihood of a cyber-attack (such as intent), Consequence Prioritization is primarily concerned with the *impact* of a potential adverse Event. Boundary Conditions should be agreed upon by all party members before generating potential Events.

Objective:

- **Adversarial viewpoint vs. entity viewpoint**
    - Adversaries will determine the degree and type of impact or damage (physical, financial, reputation, etc.) from a cyber-attack when establishing their objectives.
    - The entity (specifically the organization's decision-making group) knows better than anyone what level of impact their organization can withstand before such an attack becomes unbearable.
    - These two viewpoints combined create the Objective in CCE.
- **Examples**
    - **Amount of supply or firm load affected**
        - This is the amount of supply (i.e., generation capacity) loss necessary to be considered significant, which may vary from asset owner to asset owner.
    - **Cost of damage**
        - This is the amount in dollars of damage necessary to impact operations or the mission.
    - **Duration of outage**
        - This is the length of outage time necessary to impact customers and business operations.

Scope:

- **Systems to be examined**
    - Based on ownership and understanding, what relevant systems, processes, and components can be investigated?
- **Constraints or exclusions**
    - An organization may not have control or oversight over certain portions of their operations (e.g., water supply, other basic utilities). These need to be identified and can be excluded from the Scope.

- Ideally, all entity assets should be made available. In practice, however, some limitations can occur and are most often due to time, financial, or legal constraints (e.g., geographical restrictions or insufficient workforce).

Boundary Conditions:

- **Combination of Objective and Scope**
  - If the Objective is based on a specific monetary threshold of one million dollars, and the Scope includes all the transmission systems of the company, the two are simply combined to form Boundary Conditions.
  - Anything that exists in the Boundary Conditions should be clearly explained in either the Objective or Scope.
- **Example Boundary Conditions**
  - "An outage directly tied to the transmission lines, substations, or connected systems (logical or physical), from which the repair or recovery exceeds the cost of one million dollars."

Events

Next, the CCE Team should generate possible disruptive Events related to the Boundary Conditions. As mentioned previously, a disruptive Event is an end effect that would significantly inhibit an organization's ability to provide the critical services and functions deemed fundamental to their business operations or mission.

As the team works to generate these Events, the ideas should not be limited to traditional or obvious forms of cyber-attacks. It is important to consider similar events that resulted from human error, engineering failures, or natural disasters. In addition, the misuse or destruction of unique digital dependencies for critical infrastructure assets should be considered. This is done to ensure that more creative—or subtle—cyber-enabled sabotage is not overlooked.

Once a full list of Events has been generated, the CCE Team should carefully review the list to screen out any Events that cannot be achieved by cyber means. Those remaining Events are considered cyber-Events that can be partially developed for evaluation.

Developing cyber-Events

Each Event approved by the CCE Team will need to have a high-level explanation added to it. This will describe, in basic terms, how the Event could be achieved via cyber-means. This often includes mention of which systems could be leveraged to accomplish the attack. It is useful to understand the following targeting considerations during this process:

Physical Infrastructure and Interdependencies

The first category of targets to consider is physical infrastructure and interdependency areas. First consider the *physical elements* that are utilized in the performance of a defined process function. Example elements to consider within the electric sector may include generation, substation, transmission and distribution lines, control center facilities, and other components of the power system. Next, identify any interdependencies or chokepoints in the infrastructure. Specific examples include:

Infrastructure Example:

> Impacts to transmission lines near a power generation facility with intent to have multiple electric infrastructure impacts at the power delivery chokepoints. The primary resulting impact of an attack on the transmission system is larger than just an impact on one line because there will be resulting power flow imbalance across the transmission network, as well as disturbances to the underlying distribution system. Additional effects would impact power generation facilities due to the loss of a delivery path for the power produced.

Methods of affecting transmission line infrastructure could include targeting the overcurrent protection of physical assets, and then mis-operating devices to cause physical effects. Transmission substations and switchyards contain a wide variety of electrical infrastructure elements that can be mis-operated to impact the energy flow on the transmission lines. These elements may include breakers, switches, transformers, protection relays, voltage load tap change, capacitor banks, and circuit reclosers.

Interdependency Example:

> For an electric utility with assets that include gas-fired electrical power generation station(s), a "chokepoint" example would be the natural gas delivery system, most typically a pipeline infrastructure. The power generation plants are dependent on the natural gas delivery system and/or natural gas supplier (in the natural gas supply chain, this describes the natural gas producer, which can often be a company separate from the natural gas delivery/pipeline asset owner). The chokepoint could be targeted directly (delivery system or production system attack) or indirectly (attack on the asset owner of the delivery system or production system).

Horizontal Application of Technology

The second category of targets to consider is locations where technology is widely deployed, either within a system or across a geographic region. Additionally, the horizontal application of technology may refer to technology that supports a function performed by multiple organizations. Consider function-specific, widely deployed ICS technologies belonging to the same technology vendor platform, like vendor-specific implementation models of PLC's, RTU's, protection relays, meters, etc. Often, single or even multiple instances/versions of these devices may be deployed throughout a critical infrastructure business enterprise for both geographically dispersed and localized asset models.

Another aspect to consider is the increased "depth" of a technology deployment; that is, there is an incentive to develop and adopt vendor solutions that integrate new and previously deployed, legacy technologies through common programming and monitoring applications. This broad and deep functional coverage within the systems is also attractive and valuable to a potential threat actor.

Horizontal Application Example:

> An electric utility may consolidate on a specific RTU vendor to drive consistency from site to site and reduce the level of system complexity for their field personnel. If a payload targeting the common device was deployed throughout a service territory through targeting and misuse of engineering or maintenance software/procedures, the corrective actions to repair/replace the compromised hardware would be extremely time consuming, if not impossible from a workforce perspective.

From a distribution perspective, consider a smart meter worm that spreads throughout a smart meter infrastructure peer-to-peer mesh network, exploiting the common protocol and common meter firmware, and leverages the built-in capability to disconnect customer power. This creates an opportunity for an adversary to target consistency in architecture, protocols, and devices. This also provides a long deployment lifecycle for valuable exploits.

Reliance on Automation and Control Capabilities

The third category of targets to consider is made up of those which inhibit an organization's automation or control functions. Within most critical infrastructure sectors there is a desire for guaranteed reliability. To achieve highly reliable delivery of services, there needs to be a system that can detect faults or system events and automatically respond or reconfigure to continue to provide services. Within most critical infrastructure organizations, there are systems and processes that have been automated in order to provide functionality that cannot be delivered manually with the necessary real-time response to ensure system reliability and safety.

Consider the various levels of the electric sector. Power generation facilities, regardless of fuel type, rely heavily on resource inputs like automated fuel management systems, feed water systems, water cooling systems, unit control systems, voltage regulation, and a wide variety of system protection controls that prevent damage or mitigate safety risks. An adversary can target any one of these automated systems individually, or he may recognize the redundancies in place and choose to misuse or manipulate multiple systems simultaneously.

Within the electric transmission and distribution systems, there are automated components designed to detect a line fault or another physical condition that may have been caused by a downed power line or pole, and automatically isolate that line through the operation of switches, relays, or breakers. In addition, other elements within the electric system may be switched in around the fault in order to deliver power to as many customers as possible, while responding to the line event. With an understanding of the recovery process, an adversary can send false data to these automated devices to cause mis-operations or reconfigure the devices in a manner so that they will mis-operate under normal conditions. The tendency for electric utilities to use common device types and communications infrastructures can make this an attractive target for an adversary.

Electric Control Center environments contain entire systems that are designed to monitor and act both manually and automatically across a wide footprint of the electric system. This may include hundreds or thousands of substation environments, dozens of power generation facilities, and thousands of miles of transmission lines. The energy management systems (EMS) located at control centers are used to keep the system in balance; however, in the event of certain conditions, a control center operator may have to intercede by increasing generation to service load or shedding load to keep the system in a reliable state. An adversary with an understanding of this capability can target the EMS components to initiate load shed events or manipulate data in a manner that makes an operator believe certain conditions exist that would require operator actions to prevent a wider scale outage.

Automation and Control Examples:

- Natural gas pipeline station volume and/or pressure control, compressor control, and station emergency shutdown sequencing, which includes modern distributed safety systems (flame, gas, etc.)
- Any "real-time" remote monitoring and/or control of assets
- Same day modifications to natural gas receipt and delivery volumes
- Timely collection of accurate volume, gas constituent, and operational parameter data in a geographically dispersed set of system assets
- Electric utility EMS and energy load balancing systems
- Power system area balancing through Automatic Generation Control and scheduling
- Power element maintenance ticketing and electronic-tagging systems
- Use of automatic load shedding schemes within the EMS (Special Protection Schemes [SPS], Remedial Action Schemes [RAS])

Evaluate Potential High Consequence Events

Determine Severity

The Boundary Conditions established previously can be used to define the first order effects. Based upon the examples above, Table 1 shows an example of how these effects can be defined as criteria for scoring purposes. If a long list of cyber-Events needs to be reduced to make the scoring process manageable, these impact criteria can be used to quickly prioritize the list to allow the team to focus on the top items. Any criteria developed for a CCE engagement should be relevant and appropriate for the organization. The following criteria are provided as examples that have been developed by electric sector subject matter experts (SMEs).

Area Impacted: Describes whether the impact of the attack scenario is geographically localized or if it impacts the entire system. Area impacted is described as a loss of load (both firm and supply) in this example, which can be translated into several affected endpoints or accounts.

Duration: Describes the length of an outage.

Attack Breadth: Describes the extent to which a targeted technology or system is deployed, resulting in adverse operational effects. The greater the span of impacted systems, the more difficult the restoration following an adverse Event.

It should be noted that in our example, attack breadth moves beyond the number of devices impacted, since this value also considers the additional resources needed for restoration, such as additional personnel or financial expenditures. For example, following a cyber-attack targeting advanced metering infrastructure (AMI), recovery efforts may be complicated by the quantity of field devices deployed.

Additional criteria can be identified to further refine the scoring. These criteria should relate to the entity's values and primary concerns. Each should be clearly defined with thresholds that can be added to the previous criteria and used in Likert scale scoring.

Safety: Describes the potential impact on safety, including injuries requiring first aid or loss of life. For example, the power system outage resulted in health hazards or mortalities directly tied to the lack of available electric power. This value considers only the direct impacts to safety and not safety issues that stem from extended outages.

System Integrity Confidence: Describes whether restoration and recovery efforts can restore system integrity with confidence following an adverse Event (i.e., a system not operating as expected or intended, or, alternatively, malicious operation conducted by unauthorized users). One factor to consider is whether the initial attack propagates in multiple systems, therefore complicating restoration efforts. All of these may negatively impact an organization's confidence in their system.

Rather than focusing on the breadth of an attack, in some cases the system exploited may be central to the functionality of a critical service (i.e., the keep inside the castle). In these cases, an organization cannot operate the same system again because the risk of a follow-on attack is too high. In contrast, an organization may have confidence in their ability to replace impacted systems or devices and return to normal functionality and operation.

Cost (including restoration): This criterion considers the direct financial loss, including restoration costs, to the organization as a result of the failure scenario. Restoration cost is the cost to return the system to proper operation, not including any legal or other reparations as a result of the failure. It also includes secondary costs, such as purchasing replacement power in order to meet the need. For example, an organization with long term contracts will be impacted less than one with short term agreements.

It should be noted that the cost will be directly impacted by the size of an organization. That is, the cost of one cyber-Event may be evaluated as low for one utility but may be evaluated as medium for a smaller utility due to the greater "balance sheet" impact for the smaller utility.

Define Scoring Thresholds

This assessment is concerned with evaluating consequences. Once the criteria are decided upon, there needs to be a way to define the extent of their impact on the organization. The criteria are thus evaluated on a Likert scale, with values typically being none, low, medium, and high (numerical values 0, 1, 3, and 5, respectively). Referring to the criteria discussed above, the thresholds can be defined in the following manner (Table 1).

Table 1: Criteria scoring thresholds.

Criteria	None	Low	Medium	High
Area Impacted (Load or Customer Count)	Inconsequential	Loss of failure to service firm load of less than 300 MW (or) load supply loss of MSC or 2,000 MW, whichever is lower.	Loss of failure to service firm load between 301 and 1,500 MW (or) load supply loss of between 2,000 MW (or MSC, whichever is lower) and 3,000 MW	Loss of failure to service firm load greater than 1,500 MW (or) load supply loss of greater than 3,000 MW
Duration	Inconsequential	Return of all service in less than 1 day (inability to serve firm load) (or) supply outage for less than 1 week	Return to service in between 1 to 5 days (inability to serve firm load) (or) supply outage from 1 week to 1 month	Return to service in greater than or equal to 5 days (inability to serve firm load) (or) supply outage for greater than 1 month
Attack Breadth	Inconsequential	Elements of the system are vulnerable to an exploit that is actively being attacked and causing operational effects, but recovery is possible using immediately available resources. These events are covered within the utility's recovery plan.	Multiple system elements have the potential to be or have been successfully attacked causing operational effects. Recovery is possible but requires additional resources (i.e., time, personnel) not immediately available.	Many system elements have been successfully attacked causing operational effects. Restoration is complicated by the dispersed deployment of devices or scale. Timeline for recovery is unknown.
Safety	Inconsequential	Low but definite risk to safety, but only within the boundaries of "onsite."	There is a definite risk to safety "offsite," beyond the boundary of the fence.	There is a definite risk to safety that may include loss of life for one or multiple people, onsite or offsite.

System Integrity –Asset Owner Confidence	Inconsequential	Asset Owner has ability to restore and is confident in restoration integrity.	Asset Owner has knowledge to restore but does not have the resources (financial, time, personnel, etc.) to restore confidence in the system.	Asset Owner has ability to restore but is not confident of restoration integrity.
Cost	Inconsequential	The cost is significant, but well within the availability of an organization to recover from.	There is significant cost for recovery, and it will require multiple years for financial (balance sheet) recovery.	The cost triggers a liquidity crisis and potential result in the bankruptcy of the organization.

Determine Weighting Coefficients

The equation below is provided for calculating the scored impact points for each cyber-Event using the previously determined values.

$$Scored\ Impact\ Points = \alpha(Area\ Impacted) + \beta(Duration) + \gamma(Attack\ Breadth) + \delta(System\ Integrity) + \varepsilon(Safety) + \zeta(Cost)$$

Notice the weighting coefficient values (α, β, γ, δ, ε, and ζ) were determined by engineering and electric sector SMEs. However, these values can and should be altered to reflect the priorities of the subject organization. Typically, these weights are scaled 1-3, with 3 being reserved for the entity's primary concerns or values. For example, if an organization believes their primary concern is safety, then the value of ε can be increased so that ε has a value of 3.

In this example, the group agreed upon the following values for each weighting coefficient.

$$\alpha = 3 \qquad \gamma = 3 \qquad \varepsilon = 2$$

$$\beta = 3 \qquad \delta = 2 \qquad \zeta = 1$$

Finalize Severity Scoring Matrix

To accommodate scoring by the CCE Team, an HCE Severity Scoring matrix is drafted from the combination of the established criteria, defined scoring thresholds, and the weighting coefficients. Table 2 provides an example with all elements present.

Table 2: HCE Severity Scoring matrix.

Criteria	None	Low	Medium	High
Area Impacted (Load or Customer Count) $\alpha = 3$	Inconsequential	Loss of failure to service firm load of less than 300 MW (or) load supply loss of MSC or 2,000 MW, whichever is lower.	Loss of failure to service firm load between 301 and 1,500 MW (or) load supply loss of between 2,000 MW (or MSC, whichever is lower) and 3,000 MW	Loss of failure to service firm load greater than 1,500 MW (or) load supply loss of greater than 3,000 MW
Duration $\beta = 3$	Inconsequential	Return of all service in less than 1 day (inability to serve firm load) (or) supply outage for less than 1 week	Return to service in between 1 to 5 days (inability to serve firm load) (or) supply outage from 1 week to 1 month	Return to service in greater than or equal to 5 days (inability to serve firm load) (or) supply outage for greater than 1 month
Attack Breadth $\gamma = 3$	Inconsequential	Elements of the system are vulnerable to an exploit that is actively being attacked and causing operational effects, but recovery is possible using immediately available resources. These events are covered within the utility's recovery plan.	Multiple system elements have the potential to be or have been successfully attacked causing operational effects. Recovery is possible but requires additional resources (i.e., time, personnel) not immediately available.	Many system elements have been successfully attacked causing operational effects. Restoration is complicated by the dispersed deployment of devices or scale. Timeline for recovery is unknown.

System Integrity—Asset Owner Confidence $\delta = 2$	Inconsequential	Asset Owner has ability to restore and is confident in restoration integrity.	Asset Owner has knowledge to restore but does not have the resources (financial, time, personnel, etc.) to restore confidence in the system.	Asset Owner has ability to restore but is not confident of restoration integrity.
Safety $\varepsilon = 2$	Inconsequential	Low but definite risk to safety, but only within the boundaries of "onsite."	There is a definite risk to safety "offsite." Beyond the boundary of the fence.	There is a definite risk to safety that may include loss of life for one or multiple people, onsite or offsite.
Cost $\zeta = 1$	Inconsequential	The cost is significant, but well within the availability of an organization to recover from.	There is significant cost for recovery, and it will require multiple years for financial (balance sheet) recovery.	The cost triggers a liquidity crisis and potential result in the bankruptcy of the organization.

The combination of the weighting coefficients and the severity threshold values will depend on each organization. For this matrix, the maximum number of impact points is 70. The total number of impact points is determined by multiplying each weighting coefficient by the highest score possible per criteria, and then adding the results together. The following equation demonstrates how the HCE Severity Score is calculated:

$$\begin{aligned} &Scored\ Impact\ Points \\ &= \alpha(Area\ Impacted) + \beta(Duration) + \gamma(Attack\ Breadth) \\ &+ \delta(System\ Integrity) + \varepsilon(Safety) + \zeta(Cost) \end{aligned}$$

$$Maximum\ Impact\ Points^{b} = \alpha(5) + \beta(5) + \gamma(5) + \delta(5) + \varepsilon(5) + \zeta(5)$$

$$HCE\ Severity\ Score = \left(\frac{Scored\ Impact\ Points}{Maximum\ Impact\ Points}\right) * 100$$

[b] Note that not all organizations will assign the value of "5" to "High." As such, there is the potential the value for "Maximum Impact Points" will vary from organization to organization based not only on how many criteria are chosen, but also on the values they assign to their scoring definitions.

Scoring Example

As an example of the scoring process, the following HCE has been assessed and scored. The reasoning and results are shown in Table 3. The CCE Team consulted with SMEs in order to assess the impact of this Event. It should be noted that the cyber-Event scored describes a system failure rather than the outcome of a cyber-attack.

Example cyber-Event:

At the commissioning of an unspecified plant, a power interruption resulted in a loss of the control system. The plant had three combustion turbines (375 MW) and planned the construction of a 178 MW steam turbine to allow the plant to operate in combined cycle mode. As a result of the loss of power and resulting loss of the DCS, the auxiliary oil pump did not start after the trip. An emergency pump also did not start after the trip, and all lube oil was lost during roll down. The damage to the steam turbine was extensive and included damage to the bearings, the rotor, the inter-stage seals and blade, which resulted in a loss of $12 million in repairs and $30 million dollars in lost income.[i]

Table 3: HCE Severity Scoring example.

Criteria	None	Low	Medium	High
Area Impacted $\alpha = 3$		1 – While the cyber-Event does not describe the area impacted, the CCE Team assessed this cyber-Event as low due to the ability of the utility to serve load via alternative means.		
Duration $\beta = 3$				5 – The CCE Team believes that the resulting outage took more than 1 month to recover, given the amount of time required for the construction of the steam turbine.
Attack Breadth $\gamma = 3$			3– As described, the CCE Team believed that multiple systems could have been impacted (i.e., balance of plant [BOP] system, safety systems). Additionally, the impact could be applied to other facilities of the utility.	

System Integrity Confidence $\delta = 2$			3- While there is limited information, this cyber-Event would force the management of a utility to operate under the premise that their system integrity has been compromised (at least until a full cyber-forensics assessment can be conducted).	
Safety $\varepsilon = 2$		1 – There is a potential for a safety risk to onsite personnel.		
Cost $\zeta = 1$			3 – The cyber-Event describes a financial loss of $42 million. The CCE Team believed that this loss is significant, and it will require multiple years for financial recovery.	

Using the scoresheet above, the HCE Severity Score was calculated:

Recall:

$$\alpha(Area\ Impacted) + \beta(Duration) + \gamma(Attack\ Breadth) + \delta(System\ Integrity) + \varepsilon(Safety) + \zeta(Cost)$$

so

$$Scored\ Impact\ Points = 3(1) + 3(5) + 3(3) + 2(3) + 2(1) + 1(3) = 38$$

and

$$Maximum\ Impact\ Points = 3(5) + 3(5) + 3(5) + 2(5) + 2(5) + 1(5) = 70$$

thus

$$HCE\ Severity\ Score = \left(\frac{Scored\ Impact\ Points}{Maximum\ Impact\ Points}\right) * 100 = \left(\frac{38}{70}\right) * 100 = 54\%$$

It is important that all original documentation, including rationale, containing HCE Severity Scores be retained for future reference. Key decisions made by the CCE Team should also be documented and retained for future reference.

Scoring Lessons Learned

Limited Information

In evaluating various cyber-Events, the CCE Team may find they are unable to answer every question for every cyber-Event due to limited information. In these cases, some cyber-Events may be evaluated as less significant, due to their lower HCE Severity Scores. In order to compare these values against the others in the sample set, all scores should first be converted to percentages before being converted to percentiles.

Included in Table 4 is a description of how the HCE Severity Score can be adjusted in the event of imperfect information. Note that the maximum impact points will change as the CCE Team alters the weighting criteria. Using the values defined above in the example, the CCE Team may evaluate each scenario against a total of 70 potential impact points, with the most significant cyber-Events receiving higher scores. In cases where limited information required the elimination of a primary criterion (in this case duration, attack breadth, or area impacted), the total number of possible impact points decreases to 55.

For clarity, the second column was included to illustrate the elimination of some criteria (attack breadth, system integrity, or cost) for the example cyber-Event in this document. For each case, a percentage score was also calculated. While this method allows organizations to calculate HCE Severity Scores in limited information situations, it should be noted that eliminating criteria also decreases the validity of the HCE Severity Score for a given scenario.

Table 4: Example of readjusting scores based on imperfect information.

	Maximum Impact Points	**Scored Impact Points**	**HCE Severity Score**
No Criteria Eliminated	70	38	54%
One Primary Criterion Eliminated (i.e. Attack Breadth)	55	29	53%
One Secondary Criterion Eliminated (i.e. System Integrity)	60	32	53%
One Tertiary Criterion Eliminated (i.e. Cost)	65	35	54%

After calculating the HCE Severity Scores, identify the top HCE for further evaluation. If multiple HCEs are identified, some cyber-Events can be eliminated based on a predetermined threshold, as depicted in Figure 2.

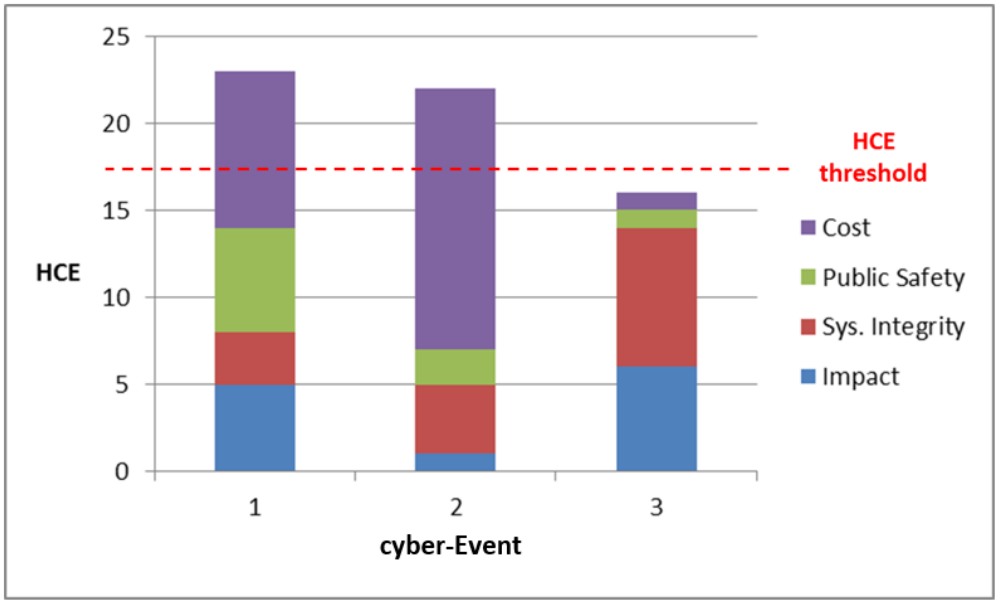

Figure 2: Example cyber-Events scored against a predetermined threshold.

Inconsistent Scoring

Each potential HCE may receive different scores from the various participants on the CCE Team. In this case, the team will need to look at the inconsistent scoring and hold a group conversation to discuss outliers to better understand the rationale for the scores given. This may cause cyber-Events to be scored and then revaluated. The team will need to decide how to incorporate these scoring changes and new rationale into the composite score. Regardless of the method chosen to combine the scores (i.e., median, average, most likely, point adjustment), care must be exercised to avoid inflating or deflating a potential cyber-Events final HCE Severity Score.

As stated previously, it is important that all original documentation, including rationale, concerning cyber-Event scoring be retained for future reference. This includes any actions taken by the CCE Team to handle scoring variance.

Revisiting Threshold Definitions and Weighting

After scoring multiple cyber-Events, the team may determine that all the scores are too similar, or that certain criteria are not given enough weight or do not provide value to the scoring process. When these situations arise, it is prudent to consider redefining, eliminating, or re-weighting the criteria to ensure that the process is functional. The CCE Team should discuss and document all changes and the rationale for those decisions.

The key takeaway is that the scoring process will likely encounter some difficulties; however, taking careful steps to correct these issues—while maintaining a group consensus—will be valuable time spent as the CCE Engagement progresses.

Validating Prioritized HCEs

After scoring is complete, the CCE Team will have identified the HCEs that are of greatest impact to the organization. This list should be prepared and presented to the entity's decision makers. This is done to validate that they agree with the group's findings and are willing to commit time and resources to the remaining CCE phases. This buy-in from the top is essential to avoid internal barriers or delays while accessing information, people, equipment, and processes necessary to conduct the engagement.

See Idaho National Laboratory's document titled "CCE Case Study: Ukraine Substation Power Outage" (INL-EXT-20-58092) for more Phase 1 examples on brainstorming Objective, Scope, Boundary Conditions, Events, cyber-Events, and criteria. The Ukrainian case study also demonstrates HCE scoring, validation, and prioritization.

---

[i] Wallace Ebner, "Strategies for the Prevention of Turbine Lube Oil System Failures," in *Proceedings of the ASME 2013 Power Conference*, July 29-August 1, 2013, Boston, MA.

CCE Phase 2: System-of-Systems Analysis

Consequence-driven Cyber-informed Engineering

Idaho National Laboratory

CYBERCORE integration center

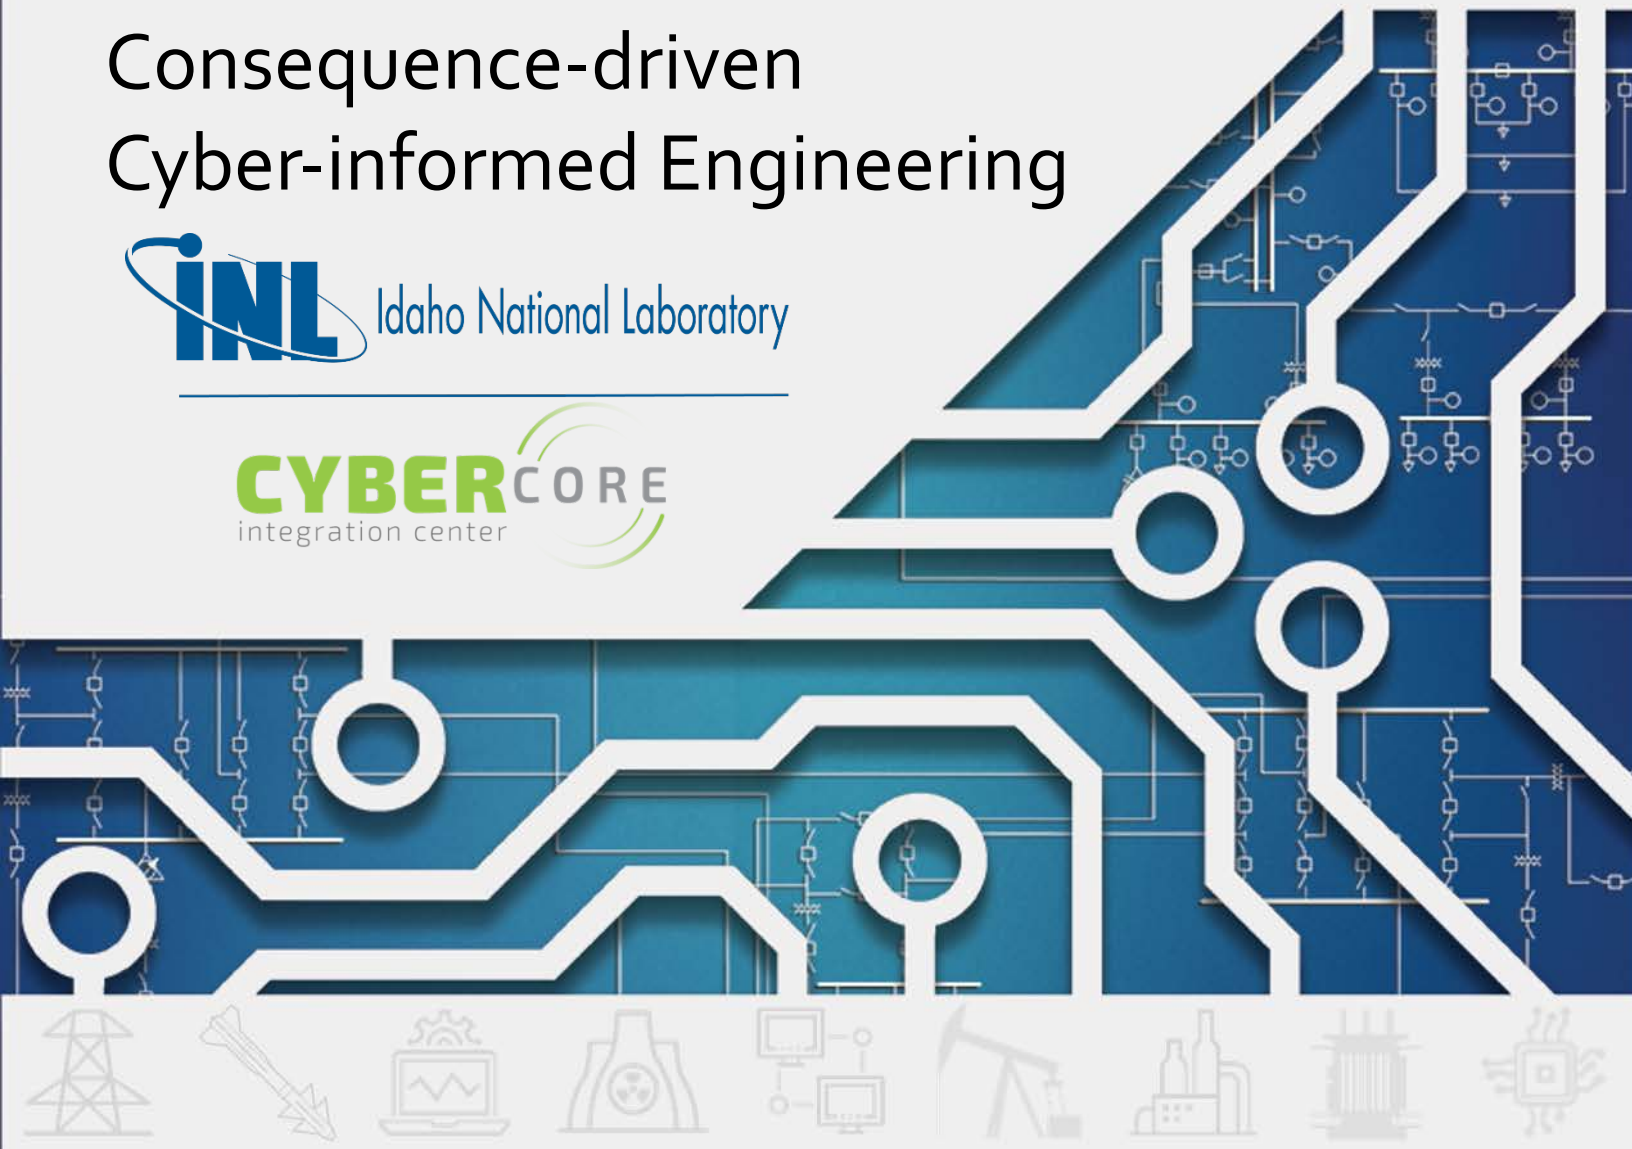

Prepared By: Doug Buddenbohm and Sarah G. Freeman
Cybercore Integration Center
Idaho National Laboratory

May 5, 2020

INL-EXT-20-58093

CCE Phase 2: System-of-Systems Analysis

Consequence-driven Cyber-informed Engineering

Doug Buddenbohm
Lead Author

Sarah G. Freeman
Co-Author

Idaho National Laboratory
Cybercore Integration Center
Idaho Falls, Idaho 83415

http://www.inl.gov

Prepared for the
U.S. Department of Energy
Office of National & Homeland Security
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517

Page intentionally left blank

CCE Phase 2: System-of-Systems Analysis

Introduction

During Phase 1, Consequence Prioritization, the CCE Team identified High Consequence Events (HCEs) that can be accomplished through cyber means to impact critical functions, services, and processes. During Phase 2, System-of-Systems Analysis (SoS Analysis), the CCE Team will conduct a systematic review and analysis of information related to the equipment, systems, processes, operations, maintenance, testing, and procurement practices based on the HCEs identified in Phase 1.

The data collected in Phase 2 will serve as the initial input for Phase 3. The SoS Analysis efforts will culminate with the System Description output, designed to summarize the information collected. The System Description functionally describes all aspects of the HCE; as such, it is exceptionally important to consider how the CCE Team will protect this data—*before* it is collected.

During Phase 2, the CCE Team focuses on collecting, organizing, reviewing, and summarizing the necessary information to fully understand the system(s) affected by the potential HCE identified in Phase 1. It is important to consider how various technologies are used within the system, what and where necessary information exchanges occur. For example, the generation site of a utility produces data that must be shared with the Energy Management System (EMS), as well as the Independent Service Operator (ISO), for balancing load. However, the specific design of that information exchange, and even the shared data, may vary from utility to utility.

At times, the operation of an organization may rely on traditional information technology (IT), as well as subcontractors, vendors, and suppliers that reside outside of the organization. SoS Analysis should be inclusive, considering all the entities, architectures, networks, and technologies relevant to an organization's critical functions or roles, regardless of location. The System Description for each HCE is the input for Phase 3, Consequence-based Targeting. A high-level overview of this phase can be found in the Phase 2 process chart on the next page (see Figure 1).

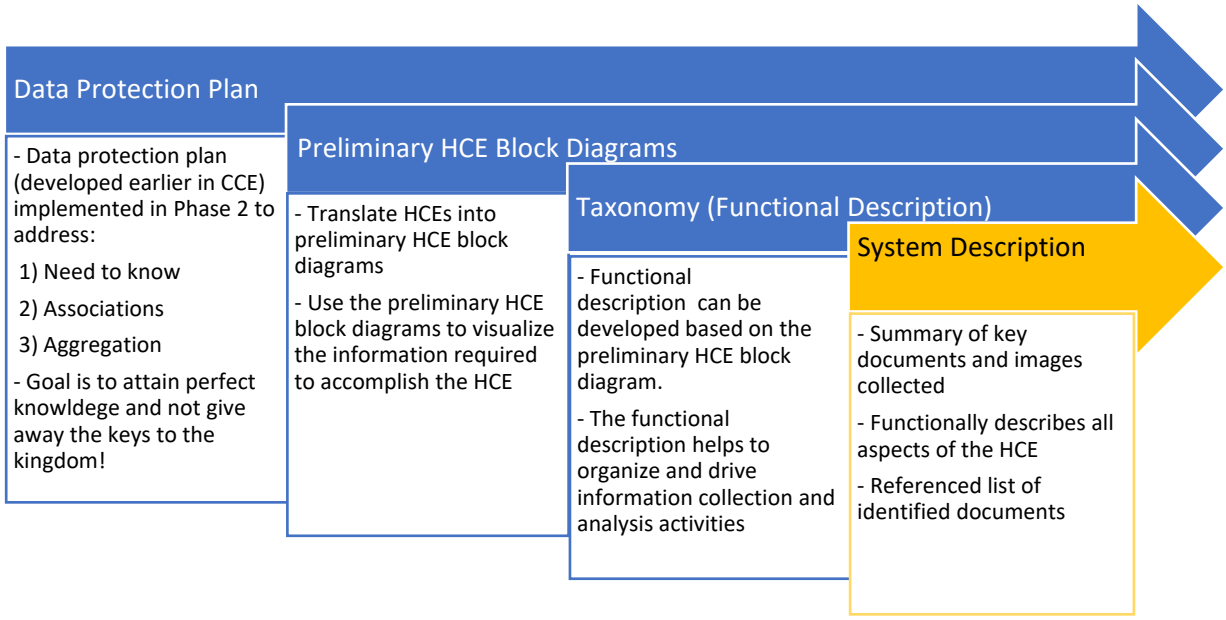

Figure 1. CCE Phase 2 process chart.

Data Protection Plan

Phase 2 collects key information that an adversary could use as a roadmap to target a system and its most important business functions. It is crucial to put in place a data protection plan to protect an organization's data. Don't give away the keys to the kingdom! The aggregation of data and documentation in Phase 2 can give the adversary full inside knowledge/access to key systems. Initially, adversaries do not fully understand the targets they have chosen—even if they have a general idea, there are still large knowledge gaps. Only an organization knows in detail how a process works, who is involved in each function, what third parties are involved, and how equipment and systems are implemented. This insider's advantage is known as **perfect knowledge**. If perfect knowledge data is not properly protected, it gives the adversary an advantage and possibly the knowledge required to successfully target key systems.

Ensure that a data protection plan is developed, properly implemented, and practiced. Equally as important, be sure the entire CCE Team understands their responsibility in keeping this information secure (i.e., not sharing or forwarding any documents, working on sensitive items on unauthorized computers/networks, or discussing the system of systems with individuals that lack a valid business reason).

Data Classification Criteria

Data should be categorized and protected according to sensitivity. Access should be limited and based on a "need to know." Information derived from the data should also be protected and categorized, based on the potential risk of damage that could occur from unauthorized disclosure. See Figure 2 on the next page for a brief description of the three criteria that factor into data classification.

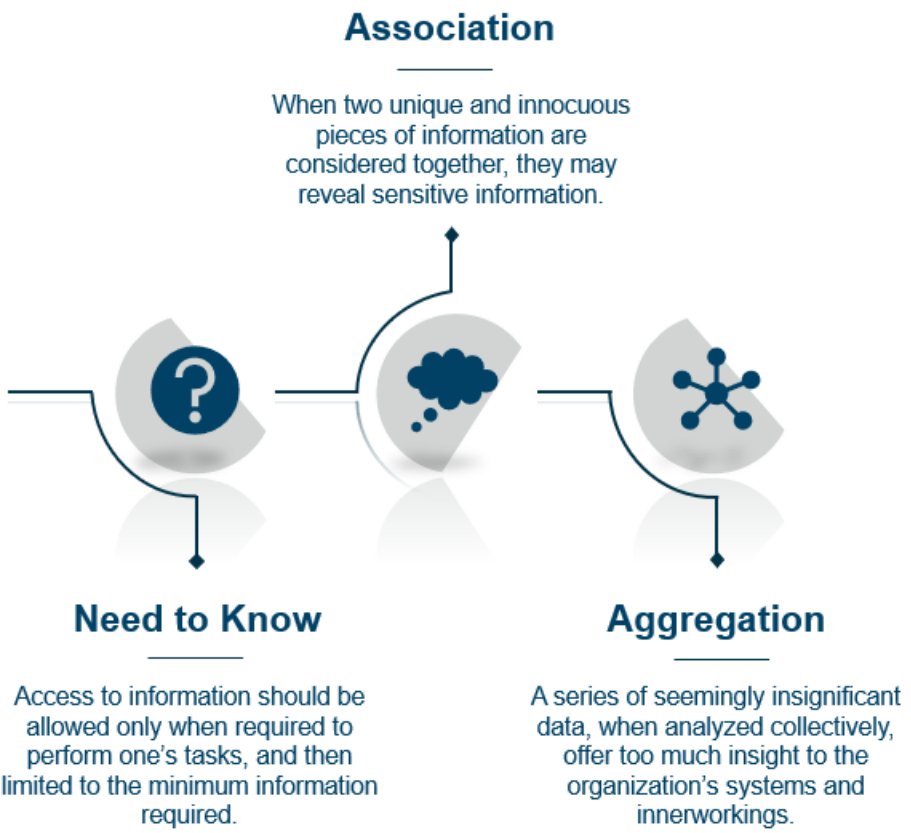

Figure 2. Data classification criteria.

Consider the following criteria that factor into data classification:

a. **Need to Know**: This is the fundamental security principle in safeguarding information. Requiring a need to know for data access ensures that such information is available only to those persons with appropriate managerial approval who have met clearly identified requirements. For example, a third-party vendor and the CISO should have different levels of access because they require different levels of need to know to accomplish their tasks.

b. **Aggregation**: Individually insignificant or apparently unimportant items or information that, when combined, reveal system details, objectives, requirements, plans, or other sensitive aspects of an organization's business mission. The disclosure of such information would provide insight into sensitive or mission critical activities, capabilities, vulnerabilities, or methods. Information amassed or collected in one location should be protected.

c. **Association**: The significance of information often depends upon its context. Therefore, when two unique and innocuous pieces of information are considered together, they may reveal sensitive information. For example, consider two unique facts: Siemens manufactures controllers, and a company publishes a job announcement for someone with Siemens controller experience. An adversary may be able to use this announcement to accurately draw a conclusion about the sensitive fact that the company uses Siemens controllers.

It is important that organizations recognize that creating this aggregated data may be dangerous for their organization, but not collecting these data (and ignoring the associated risks that already exist) is more dangerous.

Consider the following types of information to protect:

- Information in a storage medium that has been removed from another information system, or information that has been inadvertently stored in or transferred through an unprotected system.
- Information describing the nature, exploitation, or location of a system vulnerability, as well as the descriptions of the procedures required to remove/mitigate the vulnerability. In situations where mitigations only partially limit exploitation, the vulnerability information is still sensitive and must be protected.
- Information that could reveal, jeopardize, or compromise a device, piece of equipment, or the technology used in a system.
- Information pertaining to a system that reveals capabilities or weakness that would provide insight or motivate an adversary to develop malware or an exploit.
- Description of the design, capabilities, and functions of an information system[a] (or software developed to process that information) could reveal a method or reduce the level of effort for an adversary to achieve an objective.
- Information that reveals organizational structure, job posting specifics, and staffing levels may provide insight to an adversary.

Preliminary HCE Block Diagrams

After revisiting and developing the data protection plan for Phase 2, the CCE Team creates a simple, high-level diagram for each HCE. These preliminary HCE block diagrams help visualize the information required to accomplish the outcome. This exercise helps narrow the scope of analysis, organize the physical and functional connections between the target components and the affected systems, and minimize the volume of information collected to describe each HCE. The preliminary HCE block diagram provides a starting point for identifying what information and system accesses the adversary needs to accomplish the HCE. This information steers the data collection efforts.

Taxonomy (Functional Description)

Most of the activity in Phase 2 will involve identifying, collecting, and organizing documentation relevant to an HCE. This information is used to build a comprehensive knowledge base of key details for the SoS Analysis. The goal is to obtain perfect knowledge of the system(s) relevant to the HCE. To help organize the collection and analysis activities, a taxonomy or functional description can be developed based on the preliminary HCE block diagram. This is often best done by starting with the target components that must be affected to cause the HCE and working backwards. Considering the following:

- What systems and equipment are involved in the HCE?
- What documentation is needed to describe interconnected systems and dependencies?
- What relationships with other entities are involved?

[a] Information system refers to any telecommunications and/or computer-related equipment that is used in the acquisition, storage, manipulation, management, movement, control, display, switching, interchange transmission, or reception of voice and/or data (digital or analog), including software, firmware, and hardware.

The functional description can be represented as a hierarchical data structure or taxonomy. Using this functional description as the basis for investigation, the CCE Team will begin collecting and organizing key details. Relevant information to support this work includes details of interconnected systems and dependencies, controllers, technical manuals, diagrams, protocols, access lists, associated manufacturers, trusted relationships, contractors, suppliers, emergency procedures, and personnel. The SoS Analysis proceeds in parallel during information collection by building an understanding of the critical systems and processes.

Recall both the data collection effort and the CCE methodology are iterative. As the CCE Team identifies specific information gaps from the SoS Analysis, time is taken to adjust the detailed information collection to close these gaps. While not all-inclusive, the resulting information will build upon the preliminary HCE block diagram. This will ideally result in a body of perfect knowledge. This will benefit the organization by both identifying critical information and determining where it resides.

For example, is the critical information on internal servers or a public-facing server? To help ensure continued data collection efforts remain focused on the HCE, it may help to build out the original diagram throughout Phase 2. This helps produce diagrams with greater detail as more data is collected and aggregated. The point of Phase 2 is to be aware of all the information that an adversary would need to execute a successful attack.

System Description

In order to analyze the system to develop a targeting plan, the CCE Team must collect as much relevant information as possible and then summarize the key details to support a deeper level of knowledge of the system operations, personnel support activities, system configuration, and other aspects of the operation. To accomplish this, a **System Description** is developed that details the key information that an adversary may need to obtain access and accomplish the HCE through cyber means. This description should detail all the elements in the preliminary HCE block diagram and provide traceability to all the information collected in Phase 2, including where it resides and who has access to it. This System Description will be the output of Phase 2 and the input to Phase 3.

See Idaho National Laboratory's document titled "CCE Case Study: Ukraine Substation Power Outage" (INL-EXT-20-58092) for more Phase 2 examples on creating preliminary HCE block diagrams, taxonomies, and System Descriptions.

CCE Phase 3: Consequence-based Targeting

Consequence-driven Cyber-informed Engineering

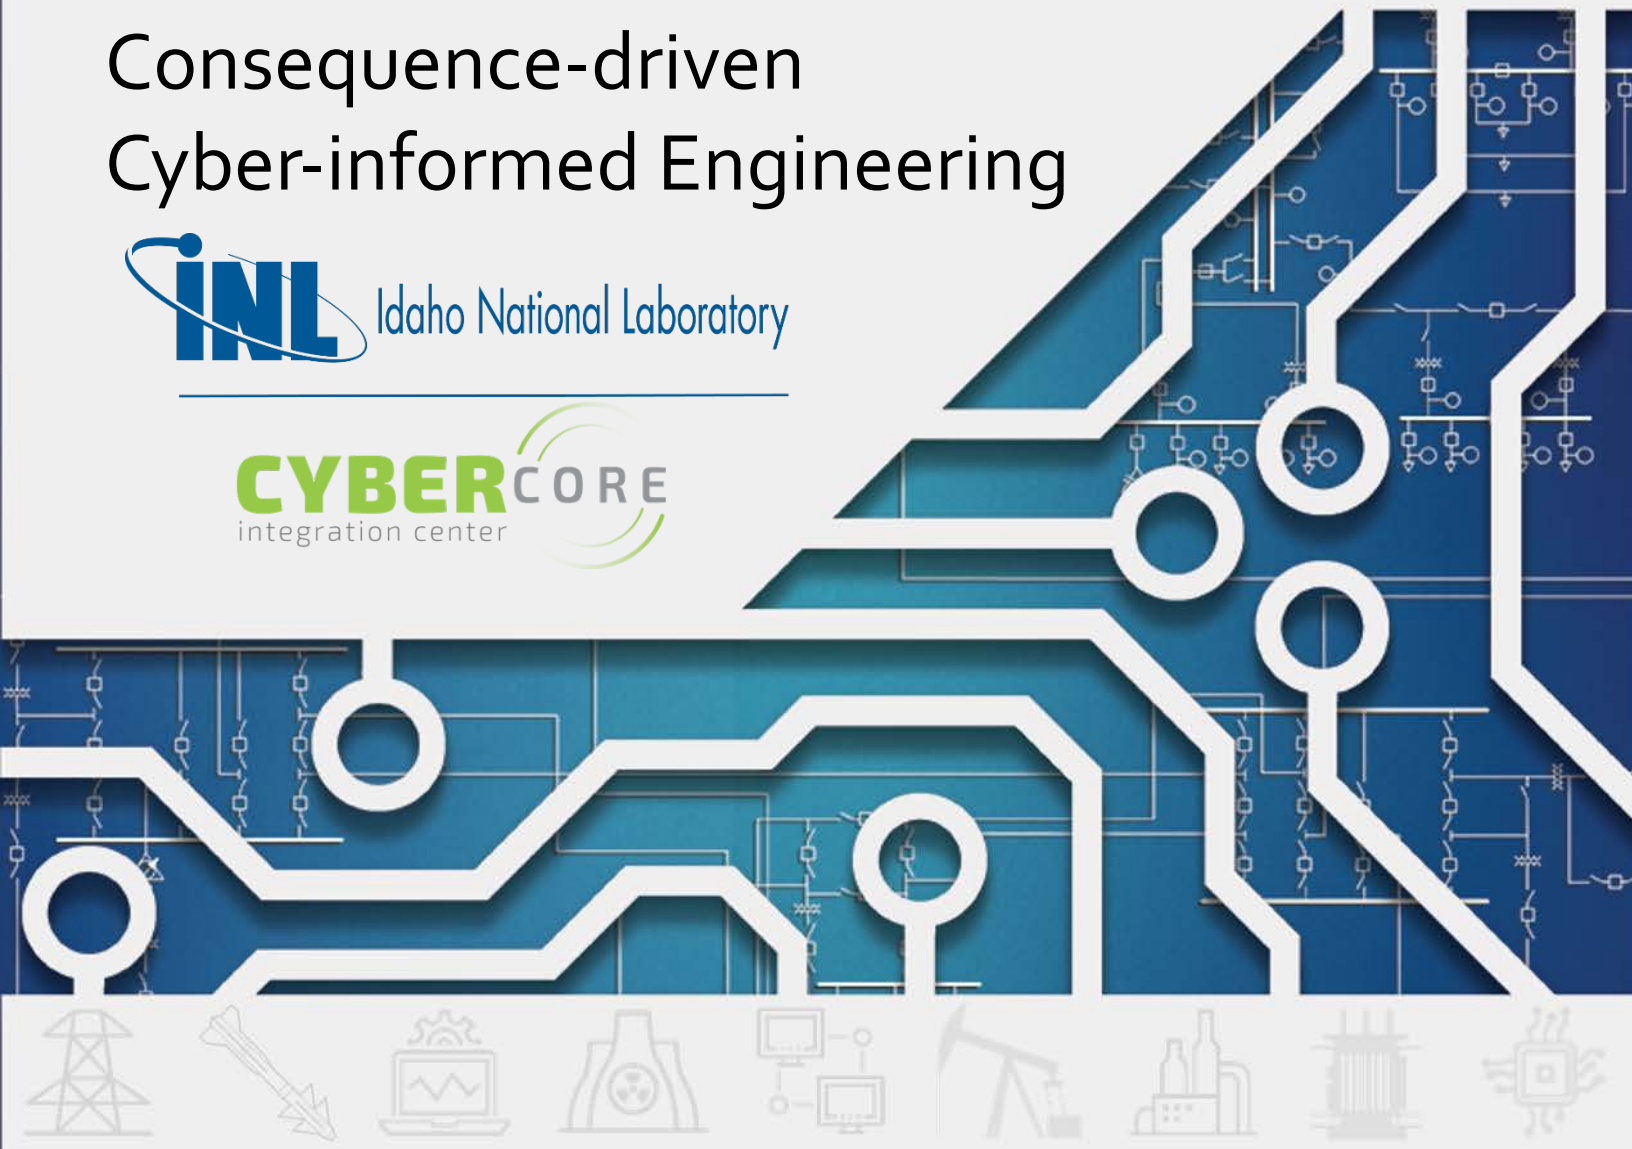

Prepared By: Stacey Cook and Sarah G. Freeman
Cybercore Integration Center
Idaho National Laboratory

May 5, 2020

INL-EXT-20-58090

CCE Phase 3: Consequence-based Targeting

Consequence-driven Cyber-informed Engineering

Stacey Cook
Lead Author

Sarah G. Freeman
Co-Author

Idaho National Laboratory
Cybercore Integration Center
Idaho Falls, Idaho 83415

http://www.inl.gov

Prepared for the
U.S. Department of Energy
Office of National & Homeland Security
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517

Page intentionally left blank

CCE Phase 3: Consequence-based Targeting

Introduction

During Phase 1, Consequence Prioritization, the CCE Team identified High Consequence Events (HCEs) that could be accomplished through cyber means to impact critical functions, services, and processes. During Phase 2, System-of-Systems Analysis, the CCE Team conducted a systematic review and analysis of information related to the equipment, systems, processes, operations, maintenance, testing, and procurement practices based on the HCEs identified in Phase 1.

A summary of the HCE-relevant information collected in Phase 2 was drafted into a System Description, which serves as the starting point for the targeting analysis performed during Phase 3, Consequence-based Targeting. The goal of Phase 3 is to develop plausible **Attack Scenarios**. The CCE Team examines the data from Phase 2 with an adversarial perspective to brainstorm different ways to achieve the HCE. The **System Targeting Description** is used to summarize and reference all the key details that are required for the Attack Scenarios.

It should be noted that the findings in Phase 3 are not all-inclusive; they represent a set of *possible* approaches, called **Technical Approaches** in CCE, that can disrupt critical systems or functions. At the same time, these identified Attack Scenarios may be limited or informed by the Boundary Conditions defined in Phase 1. The **Target Details** describe each location where manipulation or compromise occurs in an Attack Scenario to make the HCE possible. Target Details include all the technical details an adversary would need.

While Phase 2 was a data collection effort, Phase 3 is a targeting effort. In Phase 3, organizations systematically identify the necessary steps for adversary success—all from the adversary's perspective. A key component to this approach is identifying the critical information needs, targets, access, and actions required for the adversary to achieve the HCE. These **Critical Needs** are tied to accomplishing the HCE, such as the technical requirements for the payload (**Development**), or the access required to deliver a payload (**Deployment**).

Critical Needs can and will be identified outside of an entity's network boundary or direct control (vendors, suppliers, subcontractors, regulatory, or financial filings) as well as through publicly available, open-source resources found in various places. An entity's ability to identify what these Critical Needs are, where they reside, and who has access to them is a crucial step in understanding—and ultimately mitigating—risk.

For the CCE Team, the definition of critical information should extend well beyond documentation because an adversary will need to understand precisely how a process or piece of equipment functions to achieve a specific effect. To gain this type of knowledge, the adversary may need to acquire equipment, software, configuration files, or even access somewhere in the supply chain. An understanding of Critical Needs can also be used as the basis for "tripwires" that flag adversary activity related to the HCE.

Visualizing Cyber-enabled Sabotage with the CCE Kill Chain

The CCE Kill Chain (see Figure 1 on the next page) was developed to help illustrate the activities an adversary must accomplish in order to cause cyber-enabled sabotage. Assembling the cumulative knowledge, capability, and access that is needed to maliciously manipulate a system requires a long term "campaign" of iterative targeting and information collection activities. The results achieved by these efforts are required for the associated payload Development and Deployment activities. They also directly relate to the success of a cyber sabotage campaign. Therefore, if a roadblock is met in payload Development, or new information or accesses become available, all the activities in the campaign will adjust to the new requirements.

The main reason for using the CCE Kill Chain is CCE's focus on understanding (and ultimately disrupting) the requirements an adversary needs to achieve the HCE. For example, adversaries may target vendors and subcontractors through supply chain or human recruitment tactics in conjunction with a cyber campaign.[1] This is done to both obtain critical information and gain necessary access for the deployment of capabilities. A highly resourced and motivated attacker may insert corrupt components or software several layers into the supply chain. An attacker might also investigate co-opting insiders or have their own agents apply for critical positions at the target organization, a subcontractor, or a vendor.

Rather than focusing on the network and cyber hygiene details for every possible cyber access point, the CCE Kill Chain will identify areas of unverified trust in the implementation, operation, or maintenance of a targeted control system. These instances of unverified trust are sources for Critical Needs an adversary requires.

[1] One concerning example of supply chain manipulation was demonstrated during the Havex campaign in 2014. During this infection campaign, the adversary intercepted and altered update packages for ICS and auxiliary equipment. This effort directly targeted the operations of its victims by piggybacking on the update process for non-internet facing and air-gapped machines.

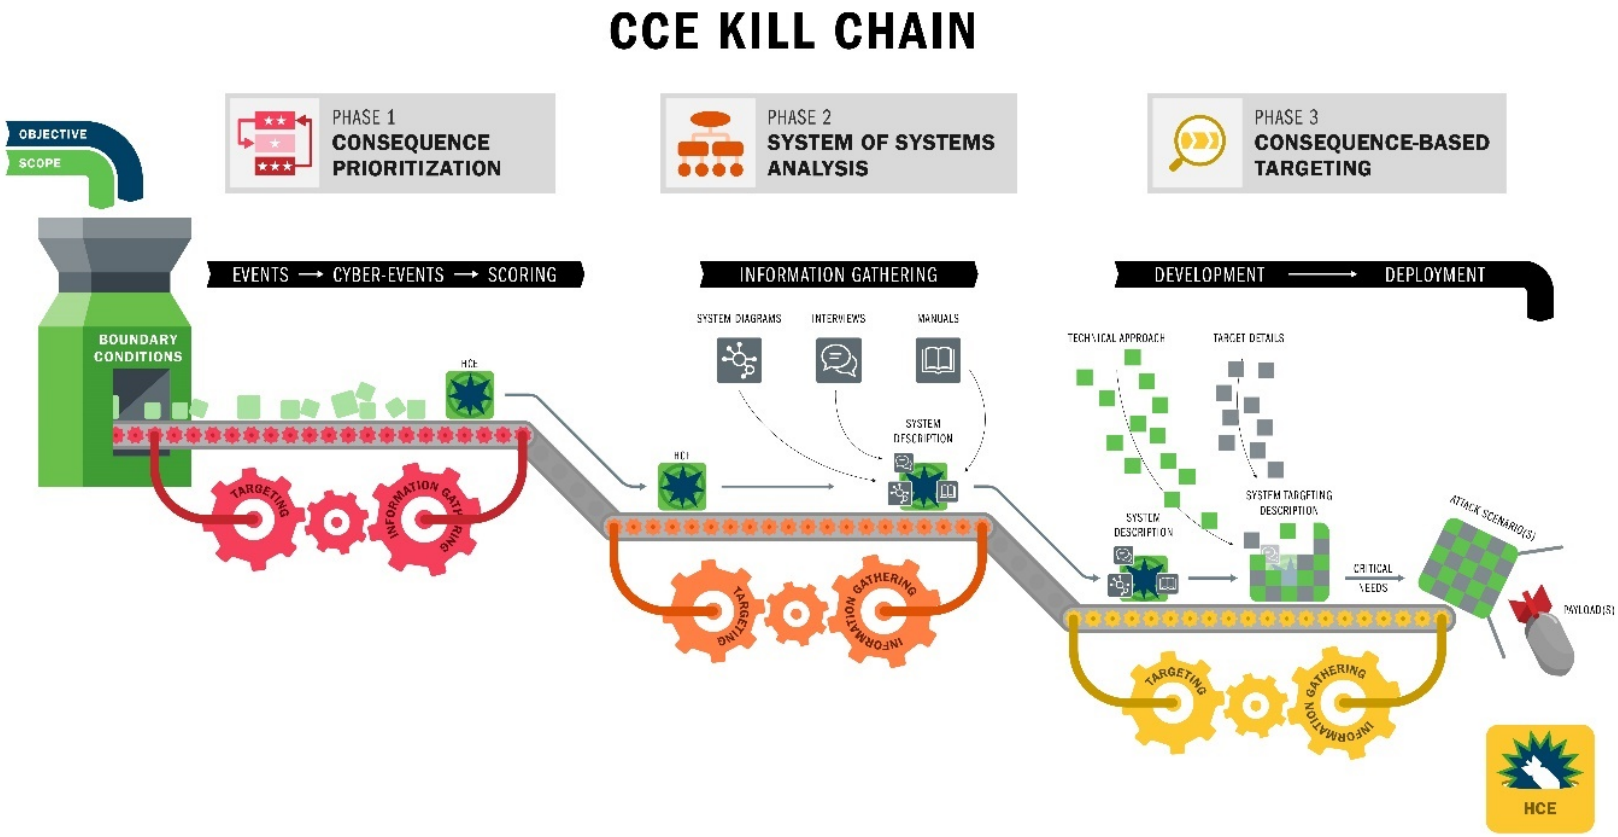

Figure 1. The CCE Kill Chain describing Phase 1 through Phase 3.

Identifying Choke Points

Think of a tree as a representation of every possible way to achieve a specific effect. The goal is to go from the tips of the branches to the roots of the tree. The branches represent all the different vectors an adversary can use to reach the goal. Even within the branches, there are many forks and possibilities. There are choke points where the branches come together—especially the trunk. These choke points are the specific locations that need to be identified by a targeting exercise.

By following the different branches of the tree down to the trunk, it is obvious there are many different pathways an adversary could take to compromise a system. Through Consequence-based Targeting exercises, these choke points an adversary must traverse are identified. These locations help narrow the areas for defenders to focus their protections to keep adversaries from reaching their goal.

Figure 2. Visual representation of choke points.

As discussed in the tree analogy, there are many different approaches an adversary could take to reach their goal of achieving the HCE. To take that concept a bit deeper, there may be multiple targets and multiple actions done to those targets, and there most likely will be multiple ways to access those targets. It is important to narrow down these numerous Attack Scenarios to the most plausible scenarios. Try to not get caught up in the countless possibilities.

Requirements for an Industrial Control System Attack

To successfully execute cyber-enabled sabotage on a critical function controlled by an industrial control system (ICS), an adversary must accomplish three basic tasks. They must develop the payload(s) required to cause the desired HCE, they must achieve access to the target ICS(s),[2] and they must get the payload(s) to the target device(s) in the ICS(s).

Like developing software, a customer provides requirements to the software developer detailing what the software must be able to do. The software developer would also need to know what kind of system the software will be running on with all the technical details. Once the software is developed, it needs to be delivered to the customer on the designated system.

While developing the criteria necessary for an adversary to cause an HCE, there are several questions that need to be answered, such as:

- "What do we have to do to achieve the HCE?"
- "Where do we have to be to achieve the HCE?"
- "How do we get to the Target(s)?"

[2] Access can be obtained through an initial access vector and/or any required network traversal.

The answers to these questions help define what the adversary needs to know, where on the ICS they need to be, what equipment or software they must access for development, and what kind of exploit they need to develop to cause the HCE.

System Targeting Description

The System Description from Phase 2 will be the groundwork that becomes the System Targeting Description in Phase 3. The System Targeting Description includes additional key details that are identified during targeting analysis and complete the summary of information required for the Attack Scenarios to cause the HCE. Adversaries have many different vectors or scenarios they could use to reach their goal. Just like the previously discussed tree analogy, there are a myriad of branches or pathways available, but the most plausible routes and choke points need to be the focus during this targeting exercise.

References are crucial in this step. Every piece of information documented in the System Targeting Description should also be referenced to the easiest accessible location. For example, if the fact that an organization uses a specific model of equipment can be found in both internal engineering documents and from the equipment vendor's website, the vendor's website should be referenced. The public-facing website is the easier path for the adversary to collect the needed critical information because it does not require breaching the company's network.

In addition, if information can be found in several locations, the reference that contains more than one piece of critical information should be cited. Again, thinking like the adversary, it is more advantageous to find several pieces of information in one location than it is several spread out. This gives the entity a good place to start with mitigations.

Consequence-based Targeting Process

The complex process of completing Phase 3 involves recording all the different Attack Scenarios possible to cause the HCE. Within these Attack Scenarios, there will be numerous targets that require specific actions and payloads based on their technical details. These targets also have numerous ways to be accessed. In building out the details required for each target, one possible Attack Scenario is synthesized. The different accesses to each target construct a possible pathway the adversary could take to reach the final goal of the HCE.

The technical details of the targets identified in the Technical Approach are described in the Target Details, which include the details of each specific element that would need to be manipulated or compromised to achieve the HCE. All information required for the different Attack Scenarios (each with a completed Technical Approach and Target Details) would be included in the System Targeting Description with complete references for each piece of information. This is an iterative process and will likely require revisiting Phase 2 activities in order to collect the pertinent information for targeting. Figure 3 illustrates the entire Phase 3 process.

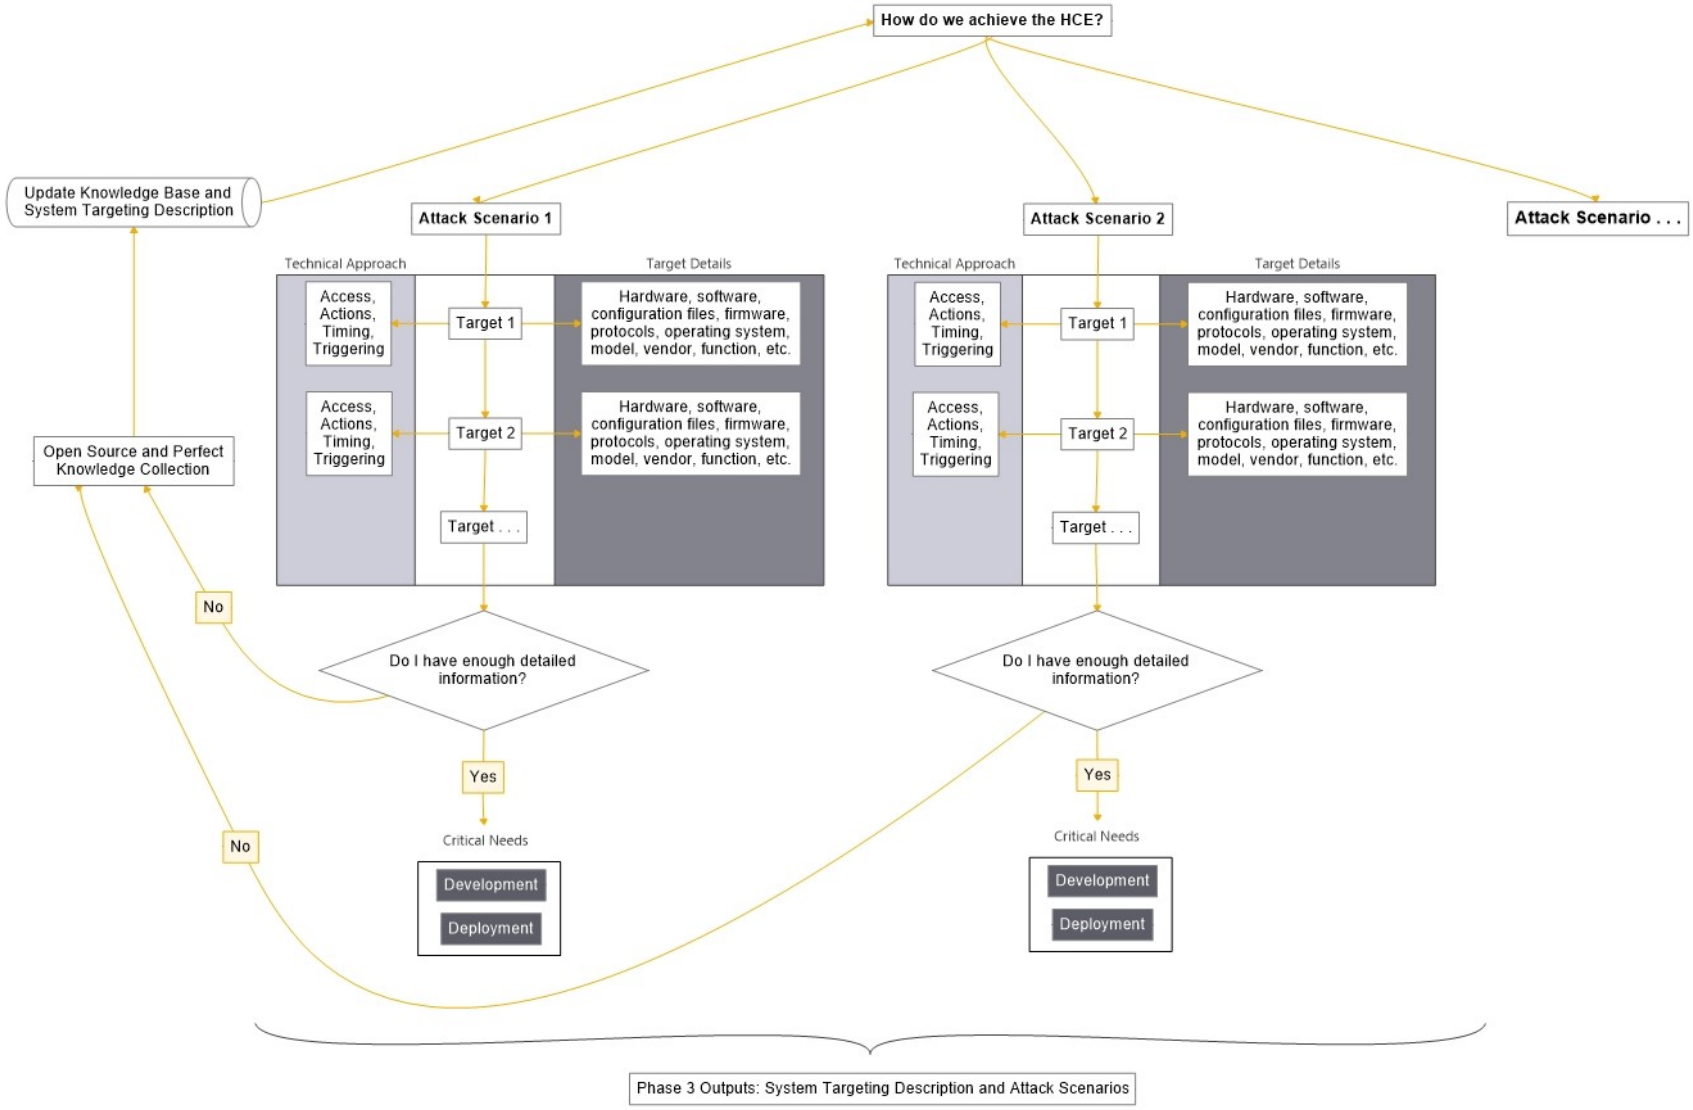

Figure 3. Phase 3 process flow chart.

Technical Approach

The Technical Approach is a detailed set of requirements. It defines a series of steps required for the exploitation of a target ICS environment to achieve a cyber-enabled sabotage effect. These individual steps will be tailored to the implementation of the targeted ICS. They will also describe how to get to each ICS element that the adversary needs to manipulate or compromise to enable the final cyber-enabled sabotage goal.

Each element in the Technical Approach will be identified by the **Target** ICS element. This specific Target might be a device component, system process, memory module, programmable chip, or logic circuit. The Target can also be non-cyber or human components of the process, like personnel with direct access to the system. Each element will define the **Access** to the target, what **Actions** need to occur at the target, and when (**Timing**) and how (**Triggering**) the payload will be triggered.

TECHNICAL APPROACH ELEMENTS

Access are the steps, movements, and actions an adversary performs to reach a target. This can occur by several pathways, but it is important for the CCE Team to choose one (either the easiest path for the adversary, or one that highlights a "blind spot" of unverified trust). It will be easy to get bogged down with all the different pathways, but it is important to focus on the most plausible pathway to discover critical choke points. In the case of extremely well-resourced adversaries, like nation-states, Access can be achieved in a variety of ways. These can include network-based, human-enabled, and supply chain methods. Keep in mind people can be either wittingly or unwittingly involved in the adversary's Technical Approach.

Actions are the conditions or steps that need to be accomplished to cause the HCE. This includes what conditions need to be met to initiate the payload, as well as what the payload actions will be once initiated. This can be anything from manipulating a control valve, opening a breaker, "spoofing" a value on a human machine interface (HMI), installing malware, escalating privileges, or even having an insider insert a USB drive into the targeted system.

Timing refers to both the *order of operations* and *sequence* of an attack, as well as the actual or "machine cycle" time in which steps must occur during an attack. These steps may be taken to avoid detection or to achieve maximum damage. The details of the timing depend on the objective of the attack.

Triggering is *how* a payload is activated, and it is always tailor-made to the process or target. Attacks are most often initiated one of two ways. They can be initiated by an attacker who is interfacing with a system or device in real time or, alternatively, by an agent operating on behalf of the attacker. These agents, also known as smart triggers, can be programmed to initiate attacks at various predetermined and defined points. The simplest triggers, and arguably not smart at all, execute a payload code based on a specific date or time input. Historically, these have been referred to as logic bombs.

TYPES OF TRIGGERS

More complex trigger designs can be used, such as conditional or process state triggers. The first trigger type, **conditional trigger**, is a trigger that initiates sections of code based on predefined and programmed requirements within the trigger's logic. For example, a trigger may "arm" the payload, but it may not proceed to the next stage of an attack until a "go" signal has been received. In this case the condition is whether the "go" signal has been received.

State triggers make up the second type of triggers. State triggers initiate the attack when the target process reaches a required state. State triggers are designed around manipulating a process. As previously highlighted, these triggers are always tailor-made to the specific component and process (they cannot be applied to a similar target without code modifications) and may require subject matter expertise to design.

Keep in mind that Triggering and Timing may occur in parallel, but this is not always the case. Adversaries may choose to trigger different parts of an attack at different times or at the same time.

Target Details

The Target Details is the section that describes the operating position(s) for a cyber-attacker; it is the "where" in the question. Where does an adversary need to be to control and execute the attack?

In some cases, it may be possible for a cyber-attacker to operate from multiple locations. For example, an adversary seeking to target electric distribution infrastructure with the effect of causing an outage may be able to attack a utility from the regional distribution management system (DMS) level. This was the case during the December 2015 attacks in Ukraine. Or, they may be able to target equipment in the substation (such as a remote terminal unit [RTU]) to be effective in causing an outage from the field device level.

If the Technical Approach is thought of as the requirements for a hacker to develop the payload, the Target Details describe the software and hardware platforms (e.g., device component, system process, memory module, programmable chip, logic circuit) that will be exploited or manipulated to implement the requirements. An adversary's terminal goal is achieved by the compromise of the items described in the Target Details via the Technical Approach and payload Deployment. The details an adversary would need about each individual Target to accomplish their goal is included in this section: make, model, software, firmware, configuration files, vendor, function, model, operating system, and protocols.

It is helpful to clearly delineate these Target Details because it may be possible to disrupt adversary activity at these nodes. In some cases, reengineered solutions, additional security measures, or improved procedures may limit the attack progression or make a target too expensive (in terms of time, money, or resources). This lessens the target's "attractiveness."

Critical Needs

Critical Needs are key pieces of data (information, equipment, or software) an adversary must acquire in order to successfully sabotage a system. By identifying this data, the CCE Team determines information that can serve as indicators or tripwires of adversary activity. These Critical Needs should be documented. Be sure to include a list of the key documents, each document's location(s), and all the personnel who have access to it. Even if a key piece of information is out of the control of the organization, that should be documented. It is important to document everything so unverified trust can be identified and addressed.

One thing to keep in mind is a Critical Need can be more than a document. Critical Needs can be pieces of equipment the adversary acquires to reverse-engineer. This allows the adversary to know exactly what will occur if the payload is triggered. It can also be crucial information that is needed. Key questions to answer:

- What you need?
- Where you can get it?
- Who has access?

Development of the Payload

Critical Needs for Development include all the information, equipment, and software needed to develop a payload. The payload is the mechanism an adversary will use to maliciously manipulate or attack a system to cause the HCE. Often, the payload is designed to target the basic functions of a system and render these functions unavailable, or maliciously use available design features.

The goal of payload Development—and its corresponding cyber-attack—is a physical effect accomplished via cyber means. In contrast to many (if not all) information technology (IT)-centric attacks, a cyber-physical attack is directed against the base functions of a system, instead of access to sensitive information. For example, adversaries targeting the wicket gate of a hydro generation station may be successful in limiting or stopping the flow of water through a dam, thereby limiting the generation output of the site.

Adversaries interested in designing payloads to sabotage physical systems need a detailed level of understanding of the target process to manipulate it for disruptive purposes. Because of the additional knowledge required, engineering design documents and other technical specifications will be a key element of the targeting process. Another exceptionally useful source of information is mechanical failure analysis or similar documentation; this information can provide valuable insight for the adversary seeking to achieve damaging or destructive attacks via cyber means.

Reid Wightman illustrated the usefulness of these design specifications when he identified a common vulnerability in a key engineering component. Wightman designed a hypothetical attack against a variable frequency drive (VFD) by rewriting the skip frequency[3] so that dangerous conditions would be obtained by the VFD during operation.[a] Wightman also noted that in many cases the skip frequency field was read/writable, allowing for potential malicious alteration by an adversary.

Deployment of the Payload

Critical Needs for Deployment of the payload include the pieces of critical information the adversary needs to deliver the payload to the intended location. Delivery of the payload often requires different accesses than those that were used during payload Development. Other considerations include the desired scale of the attack and how many systems will need to be sabotaged to achieve the HCE.

For example, if an adversary wants to affect an entire fleet of ships—and not just one ship—the Critical Needs for Deployment will be different. They will need to figure out how to deploy their payload to all the ships and not just one. This may be achievable through the supply chain. If the entire fleet relies on one common vendor for a target component, the adversary may only need to interrupt the supply chain in one location. However, if the ships used different suppliers for the target component, the Deployment may require access to the supply chain in more than just one location.

Documentation and Reporting

Each Attack Scenario should be drafted with the key collected information summarized in the System Targeting Description. This will help inform a thorough and knowledgeable presentation for the company's C-Suite. Being able to translate targeting information into their language (risk, cost, efficacy, consequence, etc.) will help them understand the risk to their business, generate their buy in, and facilitate the implementation of mitigations in Phase 4.

Outputs and Next Steps

The output of Phase 3 and input to Phase 4 are fully developed Attack Scenarios that can accomplish the HCE and a fully documented and referenced System Targeting Description. Each identified Attack Scenario will include:

- Technical Approach with the requirements for each target including the Access to the target, the Actions needed to be taken, and the Timing and Triggering of the payload.
- Target Details which describe the technical details of each target that will be exploited or manipulated in order to implement the requirements from the Technical Approach.
- Critical Needs, which describe what an adversary requires (information, access, components, software, etc.) for both payload Development and Deployment, including the "easiest" place to obtain them.

With all the Attack Scenario details and choke points compiled into a Phase 3 summary document, the CCE Team can use them in Phase 4 to articulate specific mitigations and protections to secure those choke points.

[3] A skip frequency is a designated frequency for a specific piece of equipment at which unsafe vibrations and other damage can occur.

See Idaho National Laboratory's document titled "CCE Case Study: Ukraine Substation Power Outage" (INL-EXT-20-58092) for more Phase 3 examples on defining Attack Scenarios, creating a System Targeting Description, and developing Critical Needs.

[a] Zetter, Kim. "An Easy Way for Hackers to Remotely Burn Industrial Motors." Wired Magazine, January 12, 2016, https://www.wired.com/2016/01/an-easy-way-for-hackers-to-remotely-burn-industrial-motors/.

CCE Phase 4: Mitigations and Protections

Consequence-driven Cyber-informed Engineering

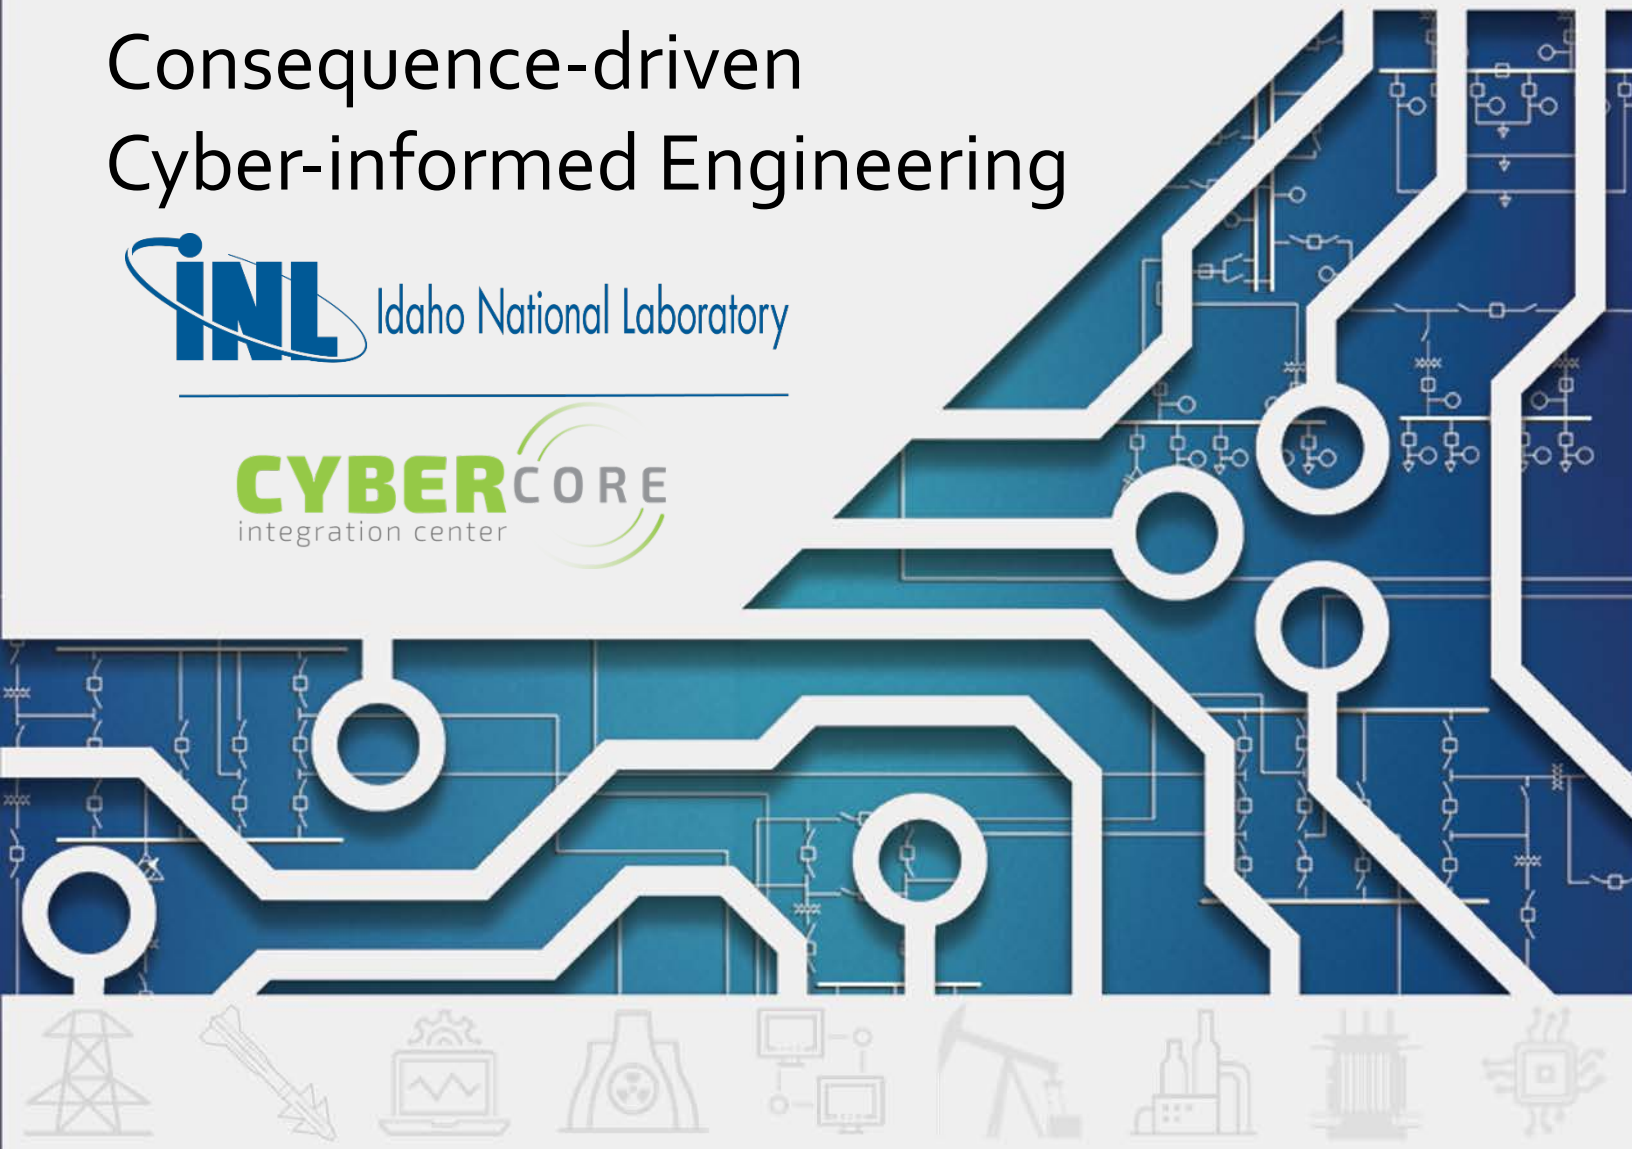

Prepared By: Theodore Miller and Sarah G. Freeman
Cybercore Integration Center
Idaho National Laboratory

May 5, 2020

INL-EXT-20-58091

CCE Phase 4: Mitigations and Protections

Consequence-driven Cyber-informed Engineering

Theodore Miller
Lead Author

Sarah G. Freeman
Co-Author

Idaho National Laboratory
Cybercore Integration Center
Idaho Falls, Idaho 83415

http://www.inl.gov

Prepared for the
U.S. Department of Energy
Office of National & Homeland Security
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517

Page intentionally left blank

CCE Phase 4: Mitigations and Protections

Introduction

During the first three phases, the CCE Team identified any instances of unverified trust in the organization's technologies, processes, and procedures, any or all of which could be used to adversely impact the system. In Phase 4, the primary goal is to remove the possibility of the end effect—that is, to develop means or mechanisms that will ensure an adversary cannot achieve their Objective (identified in Phase 1) via cyber means. Such measures are known as "protections."

In some cases, this may not be possible, or the implementation of protections may not be desirable due to other considerations. In such cases, means and mechanisms should be developed that focus on putting an organization in a better position to identify adversary activities directed against it, increasing the cost of cyber-enabled sabotage for the adversary (including making things more difficult for the adversary and attempting to lower the chances an adversary may succeed), or decreasing the recovery cost of a victim organization. These measures are known as "mitigations."

Categorizing CCE Mitigations and Protections

In CCE, mitigations and protections are categorized by their function. Some options are designed to completely stop an attack, whereas others are implemented to thwart or discourage an adversary from being successful. The mitigation and protection functions in CCE were inspired by NIST's Five Functions.[1] CCE employs the NIST framework as a guideline for categorizing mitigations and protection options identified during Phase 4. As of April 2020, NIST's website[a] listed NIST's five functions as:

1. **Identify:** Assists in developing an organizational understanding to managing cybersecurity risk to systems, people, assets, data, and capabilities.
2. **Protect:** Outlines appropriate safeguards to ensure delivery of critical infrastructure services.
3. **Detect:** Defines the appropriate activities to identify the occurrence of a cybersecurity event; enables timely discovery of cybersecurity events.
4. **Respond:** Supports the ability to contain the impact of a potential cybersecurity incident.
5. **Recover:** Identifies appropriate activities to maintain plans for resilience and to restore any capabilities or services that were impaired due to a cybersecurity incident; supports timely recovery to normal operations to reduce the impact from a cybersecurity incident.

The focus of Phase 4 in CCE is on the last four of NIST's Five Functions: Protect, Detect, Respond, and Recover. The Identify function is covered in previous phases, particularly during Phase 2.

[1] The NIST Five Functions were developed in response to the February 2013 passage of Executive Order (EO) 13636, "Improving Critical Infrastructure Cybersecurity." As part of this order, National Institute of Standards and Technology (NIST) was directed to work with stakeholders to develop a voluntary framework, based on existing standards, guidelines, and practices, for reducing cyber risks to critical infrastructure. The Department of Homeland Security has used this framework to assist critical infrastructure organizations in aligning their security goals with available resources. Additional information can be found at https://www.nist.gov/cyberframework/online-learning/five-functions.

To better address CCE's goal of protecting critical infrastructure, CCE uses its own definitions for these four functions; however, the principle behind each remains largely the same:

1. **Protect:** The ability to remove the objective of cyber-enabled sabotage (take it "off the table" for an adversary)
2. **Detect:** Enables timely discovery of adversary activities.
3. **Respond:** The ability to contain or disrupt adversary activities.
4. **Recover:** Timely restoration of critical functions and services.

Protections address the "Protect" function, while mitigations address the "Detect," "Respond," and "Recover" functions. These three do complement each other, and some mitigations may address elements of all three.

Protect

CCE places heavy emphasis on the "Protect" function, as such actions will—if implemented properly—effectively make it impossible for the adversary to cause a given HCE via cyber means.

As an extremely simple example, consider a liquid storage tank at an industrial facility. A PLC governs the pump controlling the fill level of the tank. If an overfill of this tank represents an HCE, a possible protection mechanism is the use of a separate float switch in the tank that, upon activation, will physically disconnect the pump from the power supply. By installing such a device, an adversary will not be able to cause this HCE by cyber means alone.

Detect

The "Detect" function focuses on creating a means to quickly identify adversary activity. In effect, this means identifying any attempt at cyber-enabled sabotage in progress—before the adversary can achieve their objective. It is important to note that the Detect function includes all types of adversary activities—not just network activity. This could also include a shipment of a critical components not arriving, a cyber-attack at a vendor or supplier, or unexplained behavior of critical systems.

If an organization can quickly identify adversary activity, that organization has a better chance at minimizing damages. The December 2015 Ukraine power outages provide an example of this. In at least one instance, months in advance of the actual outage, a Ukrainian power company detected an adversary in the corporate network and took corrective action. Although they were not successful in taking away adversary access to their network, the hassle this presented to the adversary may have spared the organization from an attack in the end. None of the victim companies detected any adversary activity in advance of the outage.

Respond

The "Respond" function seeks to equip the appropriate individuals with a plan regarding what to do if an Event is in progress. The response plan should help to contain, disrupt, or otherwise prevent further adversary activity.

Recover

The "Recover" function aims for the complete restoration of normal operations, including whatever actions are needed for that to occur. This may differ from returning to operation, particularly if it is possible to operate in a degraded state.

Development of Mitigations and Protections

To begin developing ideas for protection and mitigation, it may be useful to conduct a structured brainstorming session. Participants for these sessions should include both the CCE Team and individuals who have not participated in previous CCE work. Ideally, participation will include SMEs *not* previously involved in the process, as such individuals may be able to examine the situation with a greater degree of objectivity.

Consider the potential advantages of having all participants in the same room at once. If this is not possible, multiple rounds of individual input may be required. Alternatively, a Delphi Method or similar design may be implemented to enable elicitation of SME expertise remotely.

The meeting should begin with an overview of what CCE is, the CCE Team's progress, key findings, and opportunities for improvement. Once everyone is on the same page, the team may choose to present all Attack Scenarios developed in Phase 3 to get a sense of the larger picture and detect any commonalities.

Brainstorming mitigations will start with a walkthrough of each Attack Scenario. During the meeting, members of the CCE Team discuss methods of eliminating threats when possible and plan mitigation actions when it is not, or as a secondary option.

As this brainstorming begins, the CCE Team will consider opportunities to strengthen security or simplify processes in a way that reduces or removes risk. To assist this effort, the CCE Team should review each fully developed Attack Scenario from Phase 3 and the associated System Targeting Description for each HCE. As each Attack Scenario is reviewed, the team may need to confirm details to assure suggested mitigation methods will be successful and not problematic for operations.

As methods are brainstormed, it is recommended that the CCE Team diagram appropriate mitigations and protections, so the entire group has a clear and concise understanding of proposed methods. This will also help the team identify patterns from one Attack Scenario to the next, allowing for the recognition of a reliance on a repeated solution and identification of potential improvements.

Prioritization of Mitigations and Protections

There is no all-purpose method for prioritizing mitigations and protections, although it may prove beneficial to at least consider the following:

- **Type ("Protection" vs. "Mitigation")**: Protections will prevent an adversary from causing an HCE via cyber means. Mitigations cannot do so.
- **Efficacy**: Proposed protections or mitigations should be reviewed for their perceived efficacy—whether the proposed solution makes the attack not feasible, or to what degree it can reduce any negative consequences or make things more costly or challenging for an adversary. This kind of review is limited in that the efficacy of a solution ultimately relies on the specific implementation that is adopted.
- **Existing Threat Information**: Some attack scenarios may leverage techniques that have already been witnessed "in the wild" (deployed against a victim) or involve targeting of systems or components that correspond with known adversary interest. The presence of existing capability or research directed against these systems will presumably result in an increased likelihood that an adversary would first pursue these options over other scenarios.

- **Assessed Attack Difficulty**: An increase in attack difficulty will presumably correspond with a decrease in the likelihood that an adversary would first pursue such a path, due to the increase in relative cost and/or the increased need for specialized skills or knowledge. In such an evaluation, the difficulty ranking can be baselined at the level of the least challenging scenario.

The most promising solutions identified will receive subsequent attention from members of the CCE Team, along with any relevant SMEs, to develop a plan for presentation to the organization's decision makers.

Implementation of Mitigations and Protections

Limited resources, such as time, money, and available personnel, will affect how quickly a protection or mitigation can be implemented. With limited resources, it is up to each individual organization to decide which protections and/or mitigations to implement, or whether they can be implemented at all.

The point is that the decision makers will ultimately be able to make a fully informed decision. The worst case is for an organization to accept risk it doesn't know about.

The CCE Team will present recommended mitigations to the appropriate decision makers within the organization. As members of the CCE Team may not necessarily "speak the same language" as individuals who work at the C-suite level, it will likely prove beneficial for the CCE Team to consciously frame such a presentation around the concept of risk management as opposed to focusing on the technical details of a given HCE and suggested mitigation strategies. Feedback from the C-suite may be incorporated as required to develop final mitigations for implementation.

Any programmatic and design changes have the potential to introduce additional risk. The eventual pursuit of any of these changes should involve a thorough cost-benefit analysis and review after a specific and detailed implementation plan has been developed. This review will determine any potential unidentified consequences and/or risks that may be introduced with these changes.

Some suggested factors to evaluate prior to the implementation of any changes are the burden and cost of implementation and maintenance—again, the emphasis on these considerations is dependent on the organization in question.

See Idaho National Laboratory's document titled "CCE Case Study: Ukraine Substation Power Outage" (INL-EXT-20-58092) for more Phase 4 examples on developing and implementing mitigations and protections.

Outcomes

The goal of Phase 4 is to develop strategies to eliminate the possibility of an adversary achieving their objective via cyber means, or to develop strategies to detect, respond to, and/or recover from adversary activity.

CCE is intended to serve as a triage activity. In some cases, it may not be possible to prevent an adversary from achieving their objective. In such cases, defensive actions should focus on increasing the resources required of an adversary to perform cyber-attacks or decreasing the resources a victim organization will need for recovery purposes.

The CCE Team is to develop these strategies. Ideally, they will identify several options. The CCE Team will then present these strategies to the organization's senior leadership team, who will ultimately make decisions regarding implementation, and how to best use the organization's resources to manage the identified risk.

[a] National Institute of Standards and Technology (NIST). U.S. Department of Commerce. "Cybersecurity Framework: The Five Functions." Created April 12, 2018, Updated August 10, 2018. Accessed April 3, 2020. https://www.nist.gov/cyberframework/online-learning/five-functions.