RELIABILITY FIRST

RISK ASSESSMENT GUIDELINES

March 2024

I. Risk Assessment Overview

In simplest terms, risk is the possibility of loss or injury. A risk assessment is a process to identify potential hazards (threats and/or vulnerabilities) and assess what the potential consequences those hazards could cause.

Risk is normally represented in the basic formula of:

(Likelihood of Occurrence) x (Potential Harm) = Risk

There are several different types of risks that organizations assess, such as: strategic, financial, operational, compliance and regulatory. These types frequently impact each other. For instance, failure to comply with a regulatory requirement could result in an operational risk to a company that also has financial consequences.

The North American Electric Reliability Corporation (NERC) requires entities to include a risk assessment of all self-reported potential noncompliance (PNC) occurrences. The assessed risk of a PNC is one factor that impacts both the method of disposing of a violation and any associated fines. In addition, those entities with self-logging privileges must be able to adequately assess risk levels, as only minimal risk PNCs can be reported thru the self-logging process.[1] Entities with self-logging privileges may have those privileges revoked for a number of reasons, one of which is unsupported risk determinations.[2] This could include routine reporting of PNCs as minimal risk, which are later determined not to be minimal risk.

Entities need to have a documented risk assessment process. This process should be performed by a qualified subject matter expert and should consistently analyze the following four key considerations:

1. The potential threats and vulnerabilities to the impacted assets.
2. The likelihood that threats and vulnerabilities could have resulted in harm.
3. The potential harm that could have resulted.
4. The combined mitigation steps that could prevent, detect, and correct potential issues going forward.

[1] Appendix 4C of the NERC Rules of Procedure, effective 5/19/22
[2] Chapter 2 of the NERC Self-Logging Program User Guide dated 11/27/2018.

II. Qualified Subject Matter Experts

Risk is inherently uncertain which could make risk assessments difficult and imprecise. Because of the imperfect nature of most risk assessments, it is strongly recommended to have risk assessments performed by trained and experienced personnel. Everyone has some grasp of risk, but not everyone is properly trained to assess risk.

Personnel need to be trained on how to make informed and reasonable estimates. They should be well-versed in the areas of the threats and vulnerabilities and what potential harm could result from them. They should have a good grasp of the likelihood that those threats and vulnerabilities could result in potential harm along with the ability to appropriately account for any risk-reducing and risk-aggravating factors that were in place during the period of noncompliance. This all factors into devising an accurate estimate of the overall risk.

One key area of training for subject matter experts (SMEs) is the technical aspects of the particular risk. This training normally has an academic foundation in a particular area of expertise (e.g., a bachelor's or higher-level degree in a particular science), and continuing education opportunities (workshops, seminars, etc.). In addition, SMEs should have sufficient experience working with the equipment and technologies in place to understand their vulnerabilities and how to protect them from misuse or misoperations. This experience is typically derived from related employment experience.

A second key area of training that companies should consider for the risk SMEs is related to improving their capability to make reasonable estimates. While everyone regularly can provide estimates for about any question, someone that is trained to make estimates has a greater likelihood of making an accurate estimate. Three integral ways to improve a person's ability to make reasonable estimates are to train them on recognizing cognitive biases, teach them to mitigate those biases, and have them complete estimation exercises with the goal of reaching a level of acceptable answers to any estimation.

When a SME has both a solid background in understanding the systems at risk and the capability to estimate risk based on known and unknown factors, they are best qualified to provide a reasonable risk assessment of a given issue. When considering a range of issues, risk assessments can be used by management to focus their resources to appropriately improve the security,

resiliency, and reliability of their systems and the Bulk Electric System (BES) as a whole.

III. Risk Assessment Process

An entity risk assessment process should be documented to assess risk consistently and accurately. It could be used by any trained SME to consistently come to the same outcome. The process should identify the methodology that is used: qualitative and/or quantitative. If the entity utilizes both, then the process needs to identify when to use which methodology. It needs to provide adequate guidelines on the completion of such risk assessment.

A qualitative risk assessment is the quickest, easiest, and most abstract. One of its purposes is to identify those risks that may need a more detailed analysis. Another use is when there is incomplete information required to perform a more detailed analysis.

One of the most effective qualitative assessments relies on the use of a risk matrix with guidelines. These are typically two-dimensional with the potential harm along one axis and the likelihood of harm on the perpendicular axis. An example is provided below:

Potential Harm based on MW lost							
	> 5k MW	Serious	Moderate	High	High	Extreme	Extreme
	2.5k to 5k MW	High	Moderate	Moderate	High	High	Extreme
	1 to 2.5k MW	Moderate	Minimal	Moderate	Moderate	High	High
	300 to 1,000 MW	Minimal	Minimal	Minimal	Moderate	Moderate	High
	< 300 MW	Negligible	Negligible	Negligible	Minimal	Moderate	Moderate
			Remote	Unlikely	Possible	Likely	Certain
			> 1 in 10,000	1 in 1000	1 in 100	1 in 20	1 in 5
			Likelihood of occurrence based on odds				

Note the potential harm example measures MW lost, but this is not the only method of measuring potential harm. An entity could use several different categories, such as: number of BES Cyber assets impacted, number of substations impacted, some other measurable loss, or a combination of these. Likewise, the likelihood of occurrence could be based on other metrics, such as how long the issue existed.

Quantitative risk assessments involve a detailed analysis of the systems involved in a potential risk. A few methods include but are not limited to: Fault Tree Analysis, Decision Tree Analysis, and Monte Carlo Analysis. Fault Tree Analysis uses a structured diagram which identifies elements that can cause system failures. Decision Tree Analysis is a diagram that shows the implications of choosing one alternative compared to another. Monte Carlo Analysis is a technique that uses both optimistic and pessimistic estimates to determine outcomes.

Regardless of whether a qualitative, a quantitative, or some combination of both is used, all methods rely on understanding the two basic inputs to a risk assessment: potential harm and likelihood of occurrence. Many also consider mitigating factors.

IV. Assessing Potential Harm

Assessing potential harm is the process of identifying adverse impacts. Common assessments include financial or reputational losses, physical damage to people or structures, or loss of assets. Potential harm in risk assessments related to noncompliance could result in any of these but should be focused on adverse impacts to an entity's systems along with impacts to the BES.

Assessing the potential impact on assets normally begins with systems within an entity's footprint. The review logically begins with the assets that are directly impacted by a given issue. Entities must be cognizant that the potential harm may not be limited to the impacted assets. Systems and assets that are interconnected or interrelated to the impacted assets may be adversely affected by the assets directly impacted by the noncompliance. Likewise, those interconnections may extend beyond an entity's footprint to neighboring systems.

NERC provides a minimum list of factors to consider in assessing potential harm:[3]

[3] Chapter 2, Registered Entity Self-Report and Mitigation Plan, dated January 2021

1. What were the system conditions during the event? For example, did the noncompliance take place while the system was stressed, e.g., during an Energy Emergency or when other emergency or special operating procedures were in effect?
2. What is the size, nature, criticality, and location of the facilities at issue?
3. How many assets were at issue and what was the nature and function of the asset(s)?
4. What other systems, facilities, or staff are exposed to the same possible failure modes?
5. Were there any misoperations, or exceedances of system operating limits or interconnection reliability operating limits (IROL) during the course of the noncompliance?
6. Was there any potential for loss of a Protection System device, degradation or loss of a BES element, loss of a BCS or information, or providing unauthorized access to BCSs?
7. Was there potential to affect any CIP technical controls that may have impacted BCSs?

Two main areas where entities often err in assessing potential harm are: failing to appropriately scope the potential harm and inappropriately considering the likelihood of occurrence or mitigating factors when determining the potential harm.

First, failing to appropriately scope the potential harm will often result in a lower risk determination as the potential harm may be less than realized. Consider the following scenario:

> An entity has established an Electronic Security Perimeter (ESP) per CIP-005. Appropriately, they have several ports open for both system-to-system and remote desktop connections. Within the ESP there is a Protected Cyber Asset (PCA) that provides data to another system outside the ESP using system-to-system communication through a designated port range. The external system is not an Electronic Access Control or Monitoring Systems (EACMS) and does not have any access to devices within the ESP except for the PCA. Software on the PCA has not been tracked by the entity for security patches.
>
> Upon discovering that the PCA's software was not patched, the entity determined it was in violation of CIP-007 R2 and self-reported the issue.

> The entity assesses that there is minimal potential harm if the PCA is compromised because it does not control any BES Cyber Assets. The entity stops short of assessing how the PCA interacts with the rest of the systems in the ESP.

In this example, the potential harm is not limited to the PCA as it is also logically connected to all the other systems within the ESP. If the PCA is compromised, a malicious actor could then execute a pivot attack from that system and begin compromising other, more critical, systems in the ESP. This vulnerability brings the entire ESP into scope for the potential harm.

Second, when considering potential harm, entities often consider facts such as: software security tools, internal controls, or the infrequency that an adverse event will happen. While these items will reduce the overall risk, they do not reduce the potential harm. A system (a single asset or group of assets) that is vulnerable to a threat that will result in a catastrophic failure will still catastrophically fail if those mitigating factors are all circumvented.

V. Assessing Likelihood of Occurrence

Assessing likelihood of occurrence is simply determining the chance that an adverse event will occur. A few common techniques include the use of words (remote, possible, likely, etc.), the use of percentages (10%, 50%, 75%, etc.) or the use of odds (1 in 10, 1 in 100, 1 in 500, etc.). The challenge is in accurately estimating how likely a particular threat would be successful in leveraging a given vulnerability.

When applying qualitative risk assessment methodologies, assessing the likelihood of occurrence may seem to be simply an educated guess, but there is much more to it. SMEs must consider the vulnerabilities that exist to the asset, what threats exist to utilize the vulnerabilities, and then estimate the likelihood of occurrence using some scale or table. Two important perspectives need to be understood when determining the likelihood of occurrence. First, just because something occurred does not mean that it's chance of occurrence is very high (e.g., very likely, 1 in 1, or 100%). Inversely, just because it has never happened does not mean that it never could (remote, 1 in a million, 0%).

Whether it's equipment failures, cyber-attacks, or human error, the chance of failure should be reasonably estimated by a subject matter expert with the proper training and experience. They should rely on resources including, but not limited to, technical or vendor publications, threat reports, and historical data. The

National Institute of Standards and Technology has a useful interpretation of likelihood: “risk assessors assign a score (or likelihood assessment) based on available evidence, experience, and expert judgment.”[4]

For quantitative risk assessments, the likelihood of occurrence is better defined through the particular tool that is used. It is important to note that any given risk assessment tool is only as good as its underlying inputs and application.

Likelihood of occurrence can be directly affected by the existence of risk-reducing and risk-aggravating factors that were in place during the period of noncompliance. Here, SMEs should consider the inclusion of internal controls in an assessment but can only take into account those that would actually mitigate the likelihood of occurrence during the time period the vulnerability existed. Internal controls that are put in place after a risk is identified will reduce the likelihood of occurrence in the future but will not impact the risk prior to implementation. In addition, internal controls that exist that are part of the vulnerability also cannot be considered in reducing the likelihood of occurrence. As to risk-aggravating factors, SMEs might want to consider overlapping issues with other security controls (i.e., multiple holes in layered security), which could increase the likelihood of occurrence.

VI. Considering Mitigation

The consideration of mitigation steps can be performed during the risk assessment or as a separate task. Mitigation steps are those that fix the immediate issue and reduce the likelihood of recurrence by preventing, detecting, or correcting future issues (i.e., internal controls).

Internal controls can be technical, procedural or both. Technical controls are automated systems that do not need human intervention to initiate. These may require manual review of the results, but those results are automatically generated, and alerts are issued. Procedural controls are the policies, procedures, and checklists that instruct or guide personnel in performing tasks.

Preventive controls are just that: preventive. They are intended to prevent a negative event from occurring. Preventive controls are an essential part of any operation as they are proactive in nature.

For instance, consider a scenario where an employee's CIP training date is entered into an electronic record, specifically a data entry field. In this scenario the entity could establish at least two procedural and two technical internal controls: First, they should have a procedural control for a documented process about how to enter employee training records into an electronic record keeping

[4] Appendix G, NIST Special Publication 800-30, Revision 1, Guide for Conducting Risk Assessments

system. Second, they should have another procedural control by establishing a checklist that guides the performer through each step of the process, including how to enter the date. Third, they could program a technical control into the data entry screen that requires a valid date entry before submitting the record. This could include the date getting entered to be within 30 days of the current date. Lastly, the date that is entered is used for a fourth internal control where there is an alert programmed into the system to alert Human Resources that an employee's last CIP training was 12 months ago.

Detective controls are designed to find issues after an event has occurred. These are essential for two reasons: They can serve as confirmation that preventive controls are working as intended and they can flag irregularities or unforeseen events that can cause issues.

Building on the last example, an entity could establish a detective control that scans the employee records data for those employees who have not completed their CIP Training within 15 months of their previous course. While the intent of the preventive control is to prevent this occurrence, no system is infallible. For instance, the system itself could suffer a technical issue or the employee could fail to complete the assigned training but not have their access removed. The records scan could be a procedural process where someone runs a report and looks for anomalies or the entity could rely on a technical control where the system generates additional alerts.

Corrective controls fix issues after they are detected. When a PNC is discovered, it is imperative that entities correct the issue as soon as possible. In some instances, permanent corrective action may not be initially possible. In these instances, entities are expected to implement temporary mitigating efforts to reduce the risk.

Lastly, despite strong procedural and technical controls, an issue is discovered where an employee did not complete their CIP training. Often this occurs due to termination or a change in role where their access was not revoked. The corrective control in this final step is that the user’s access is revoked upon discovery.