Evaluating Cyber Risk in Engineering Environments: A Proposed Framework and Methodology

Rebekah Mohr

Evaluating Cyber Risk in Engineering Environments: A Proposed Framework and Methodology

Written by
Rebekah Mohr

March 2016

Contents

1. Problem Statement

Industrial Control Systems (ICS) show an increase in the use of advanced process control software, mobile applications, web-based infrastructure and embedded IT. This implies an ever increasing use of IT hardware and applications. While the use of 'off-the-shelf' IT has brought costs down, the increased use of IT has also added complexity and increased susceptibility to threats associated with IT Security. Engineering environments are increasingly exposed to cyber risk, and the threats are increasing in both sophistication and frequency at a continual upward trend.

Figure 1: Publically known ICS-specific vulnerabilities, exploits and malware

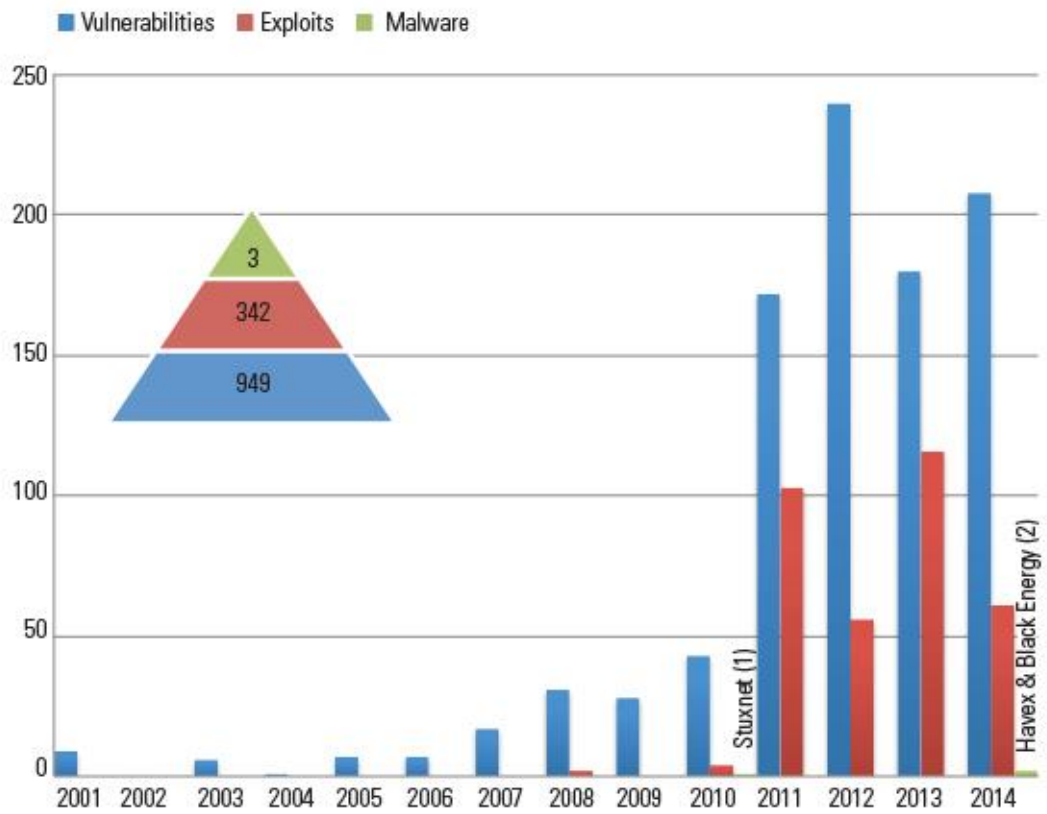

Sources: ics-cert.us-cert.gov, osvdb.org, rapid7.com, gleg.net, scadavulns.com, global data aggregated by Chris Sistrunk

The challenge facing management in Engineering organizations due to this increased trend is how to effectively assess this risk and make the appropriate investments in risk mitigation activities to manage this risk to as low as possible. Contributing to this challenge is that the methodologies and approaches available today to assess risk to operating industrial environments have been developed for traditional IT domains, and when applied to Engineering environments, are not fit for purpose as factors such as exposure, threats, and consequences are different. Moreover, the language used in IT oriented assessments differs from that used in Engineering, which can lead to decreased adoption, understanding and communication by the Engineering community. ICS are traditionally owned and managed by the Engineering community, which makes it even more difficult to communicate across boundaries between IT and Engineering. Due to the hybrid IT and Engineering nature of ICS, a hybrid competency skillset is required to secure the ICS and ensure operation is not interrupted, and because the Engineering community may not be as familiar with the threats and risk associated with IT Security, traditional IT Security Risk Assessment processes cannot be effectively leveraged. As a result, adoption of risk management is limited, leading to ineffective investment decisions and unnecessary costs.

A possible solution is to leverage commonly used Engineering risk management methodologies and incorporate IT security practices to bridge the gap between the Engineering and IT communities. This paper will explore one such methodology, the benefits and the results.

2. Abstract

The cyber threat landscape is dynamic and attacks on ICS infrastructure are increasing in both frequency and sophistication. According to ICS-CERT, 245 incidents were reported in 2014. As a result, cyber risk assessments need to be able to predict all conceivable threats which may not have already occurred in order to proactively defend against them. Instead of evaluating the likelihood of these threats based on past prevalence, they can be evaluated based on the resources and motivation which would be required to carry out an attack. The dynamic nature of this environment also necessitates consideration of seemingly endless threat scenarios. This can be limited by considering all threat scenarios based on specific variables – agent, authorization and motivation.

The Engineering community traditionally assesses process safety risk using a Bow-Tie Model and Risk Assessment Matrix (RAM), which are used to determine whether Residual Risk is As Low As Reasonably Practicable (ALARP). This phrase indicates the acceptance that risk will never be fully eliminated, and instead the goal is to determine the risk level which is considered ALARP when evaluating the cost of prevention against the potential cost of a consequence.

With some simple changes to this Risk Assessment Methodology, these tools can be used to assess cyber risks, define a Standard of cyber controls for ICS, and qualify Residual Risk in a Risk Model. These tools are critical both to communicate the risk and resulting Standard to the Engineering community, and to evaluate risk using a common language within a process environment.

After a Bow-Tie Model risk assessment has been conducted, resulting in a Risk Model, the barriers required will make up the Standard, and the depth of the controls will be defined by the determined ALARP Residual Risk. The Engineering community will be able to relate to the

resulting Standard and controls when they are translated in terms of the Bow-Tie Model, and the Residual Risk is translated in terms of the RAM.

3. Introduction

There are generally accepted methods used by the Engineering community to assess process safety risk, however this does not directly translate into assessing cyber risk within ICS because cyber threats are not predictable based on past instances due to the rapidly changing nature of the threat landscape. Moreover, both IT and Engineering knowledge is required to assess cyber risk, and both of these communities need to communicate the same way in order to effectively assess cyber risks together and put barriers in place to prevent them. A method to do this is presented below, which begins with the Engineering Bow-Tie Model, typically used to assess process safety risk. Risk is measured by running threat scenarios through the Bow-Tie Model, assessing how effectively barriers will mitigate the likelihood of a top event occurring, and then how effectively barriers will mitigate the impact of the potential consequences of a top event if it does occur.

Figure 2: Engineering Bow-Tie Model

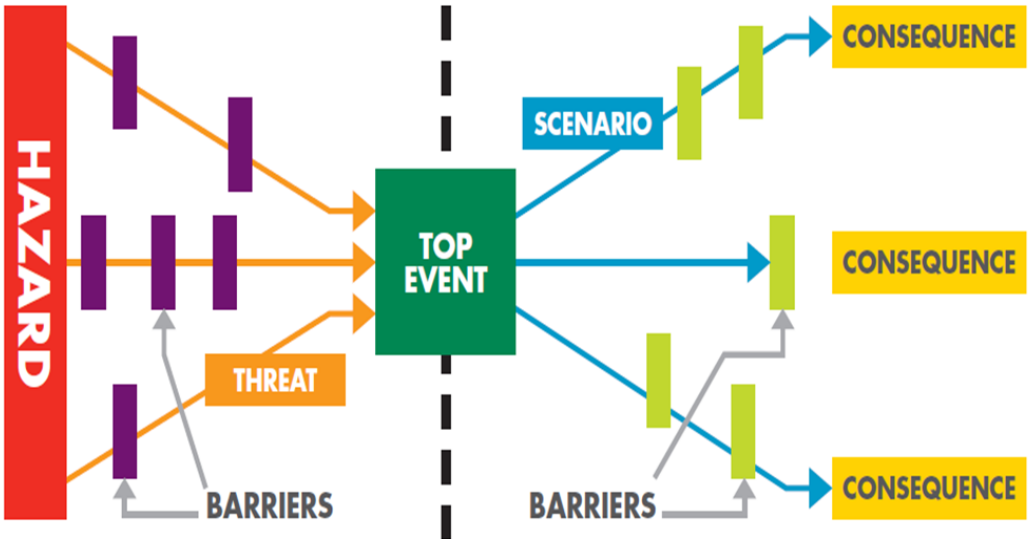

This Model evaluates all hazards which could result in a particular top event. A top event is considered a “first cause” of all possible consequences. A top event does not necessarily result in an incident or consequence. This is dependent on the barriers put in place to reduce the

likelihood of the top event occurring in the first place and to reduce the impact of the top event after it has occurred. As an example, consider driving as a hazard, and there is a threat of ice due to driving in winter conditions. A top event occurs when a vehicle loses control on the ice, and the consequence is dependent on the barriers. A barrier prior to the top event could be the use of winter tires, preventing the loss of control in the first place. A barrier after the top event could be winter driving training, which allows you to correct the vehicle and prevent any negative consequences.

This Bow-Tie Model can be translated into cyber risk by considering the current threat landscape and evaluating all possible top events. Because top events are the "first cause" of a consequence, this list is not infinite, regardless of the fact that the list of possible consequences would be infinite. Based on the finite list of top events, a Risk Model can be created which evaluates the barriers, or security controls, which are required to reduce the assessed risks to ALARP.

4. Proposed Solution

Cyber Risk Assessment Methodology for Industrial Control Systems

The first step towards obtaining Engineering community buy-in for and defining an ICS Security Standard is to compare the concepts in the cyber security industry (such as the NIST Cybersecurity Framework) to the Bow-Tie. The cyber security controls required to protect ICS Components (ICS hardware and software) can be thought of in comparison to the Bow-Tie Model by evaluating the controls that would be required to reduce the likelihood and impact of top events. While many security controls can be used to reduce both the likelihood and the impact of a top event, the controls can be categorized into Practice Areas which contain controls that typically have a function of acting as a barrier on either side of the Bow-Tie.

Figure 3: ICS Security Bow-Tie

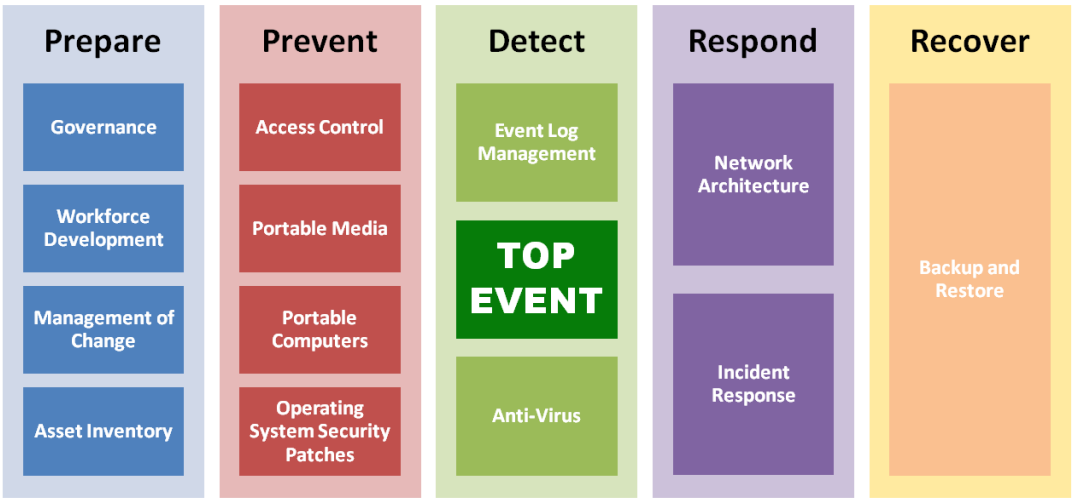

This method of communicating the required Practice Areas within the ICS Security Standard is relatable to the Engineering community, and a cyber risk assessment using a typical Engineering RAM can be used to quantitatively evaluate the Initial and Residual Risk based on the depth of the controls added or removed from each Practice Area.

Figure 4: Example Engineering Risk Assessment Matrix

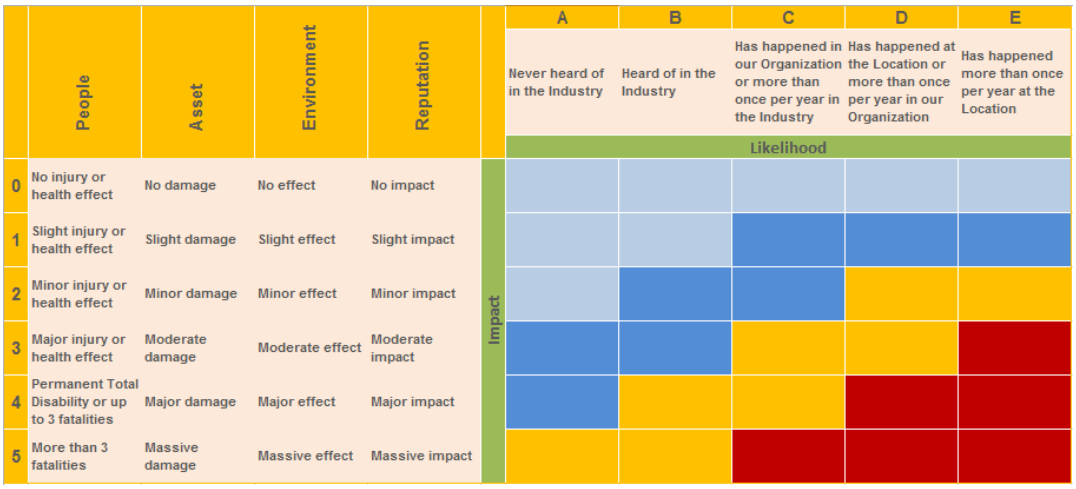

The Bow-Tie Risk Assessment methodology can then be applied to conduct a cyber risk assessment by running cyber threats through the Bow-Tie Model in the same way that hazards are evaluated. However, the threat landscape in ICS security is vast and rapidly changing. A method needs to be used to pragmatically limit the possible scenarios in order to effectively evaluate all Residual Risk. One way to do this is to reduce all possible threat scenarios to permutations of relevant variables, thereby limiting the threat scenarios to allow for Residual Risk evaluation for all possible threat scenarios. All Threat Scenarios can effectively be reduced to the following binary variables:

- Agent: The medium through which the top event materializes. If the threat is materialized through someone else that is coerced to knowingly or unknowingly exploit a threat on another's behalf, the agent is still determined by how the interaction with the target occurs.
    - Direct: The top event materializes through direct interaction with an ICS device. Example: Adversary launches an attack physically at the computer by typing commands on the keyboard or through removable media.
    - Indirect: The top event materializes through the network. Example: Adversary attacks the network.
- Authorization: How access is gained to the system in order for the actor to use the agent to compromise the process. This implies that the abuse of authorization is categorized as internal (e.g., rebooting a computer which you should not have been authorized to reboot) because the individual had the ability to carry this out using access which was not publically available.

  - Internal: The actor has some form of access into the organization which is not publicly accessible.
    Examples: Supplier, vendor, contractor, or employee
  - External: The actor does not have access into the organization, or only has access to public information/locations.
    Examples: Hackers or Anonymous without internal access/information
- Motivation: The intention behind the actions of the actor.
  - Adversarial: Individual, group, organization, or government that conducts or has the intent to conduct detrimental activities (malice).
    Example: Nation state attack
  - Non-Adversarial: Accidental or unintentional activities. Note that environmental and structural events (such as hurricanes, power failure, etc.), are not included within the definition because they are not within the ICS Security scope of control.
    Examples: Human error, equipment failure

This approach reduces the threat scenarios to 8 possible permutations.

Table 1: Threat Scenario Permutations

Agent	Authorization	Motivation
Direct	Internal	Adversarial
Direct	Internal	Non-Adversarial
Direct	External	Adversarial
Direct	External	Non-Adversarial
Indirect	Internal	Adversarial
Indirect	Internal	Non-Adversarial
Indirect	External	Adversarial
Indirect	External	Non-Adversarial

Each permutation can be tied to a threat story to make it more relatable. For example, an example of Direct-Internal-Non-Adversarial would be a contractor using removable media to inadvertently install a generic virus on ICS Component.

The type of ICS Component involved can essentially be limited to Essential or Safety Systems within an environment with Process Safety consequences or one without. This results in 3 possible Exposure permutations because an environment without Process Safety consequences would not have Safety Systems (presumably).

Table 2: Exposure Permutations

Component	Environment
Essential	With Process Safety Consequences
Safety Systems	With Process Safety Consequences
Essential	Without Process Safety Consequences

This is not to say that Non-Essential ICS Components do not play a role in evaluating the Residual Risk of threat scenarios. Non-Essential systems may be used to carry out a threat scenario, however they will likely be used as a hopping point to carry out the resulting scenario on an Essential or Safety System or to send data externally from an Essential or Safety System. Therefore, evaluating a Non-Essential system as an exposure permutation in any threat scenario does not add value because the Initial and Residual Risk would likely be higher if the exposure permutation involved an Essential or Safety System.

The next step in the risk assessment would be to evaluate each threat scenario and exposure permutation against all possible top events. Recall that top events are the “first cause” of any possible consequence. Therefore, while there could be infinite possible consequences of a threat scenario, there are a limited number of possible top events which could be evaluated against all possible threat scenarios. For example, a generic virus infection is a top event. It is considered a first cause because it would be the first event in a series of events which could lead to multiple different consequences depending on the capability of the virus.

For each threat scenario, exposure level and top event combination, the consequence needs to be determined, assuming no security control barriers and that the worst case feasible scenario

takes place. With the term “worst case feasible scenario”, it is intended that a feasible scenario is considered to ensure that unrealistic threat scenarios are not used as a basis for cyber risk management, requiring unnecessary costs to reduce the risk to ALARP. It is also intended that the worst case scenario is considered in order to ensure that the controls within each Practice Area will reduce the risk of the worst case scenario to ALARP. The worst case feasible scenario will allow for evaluation of the Initial Risk using the RAM, assuming no security controls are added as barriers.

By adding and removing individual controls from each Practice Area, an ICS Security Standard can be defined by determining the right amount of controls to limit the Residual Risk for all possible threat scenarios to ALARP.

In order to apply this cyber risk assessment methodology, some assumptions need to be made:

A. The typical RAM shown in Figure 3 evaluates the likelihood of a top event occurring based on past experience. This is applicable for non-adversarial threats, as the past accidental, structural and environmental risks are mostly similar today as they were in the past. However, the RAM does not take into consideration that the past is not a good indication of the future when it comes to adversarial threats. The ICS cyber threats and vulnerabilities have changed significantly over time, and the RAM likelihood does not reflect this. With this in mind, a separate RAM should be used to risk assess adversarial threat scenarios. This modified RAM takes into account both the motivation of the adversary, and the capability required to carry out the adversarial action. It is assumed that the likelihood is lower for an adversarial threat which requires high adversary capability (skill or resources) or for which there is low motivation. Likelihood becomes high for an adversarial threat if the adversary capability required is low and if there is high motivation.

Figure 5: Adversarial Risk Assessment Matrix

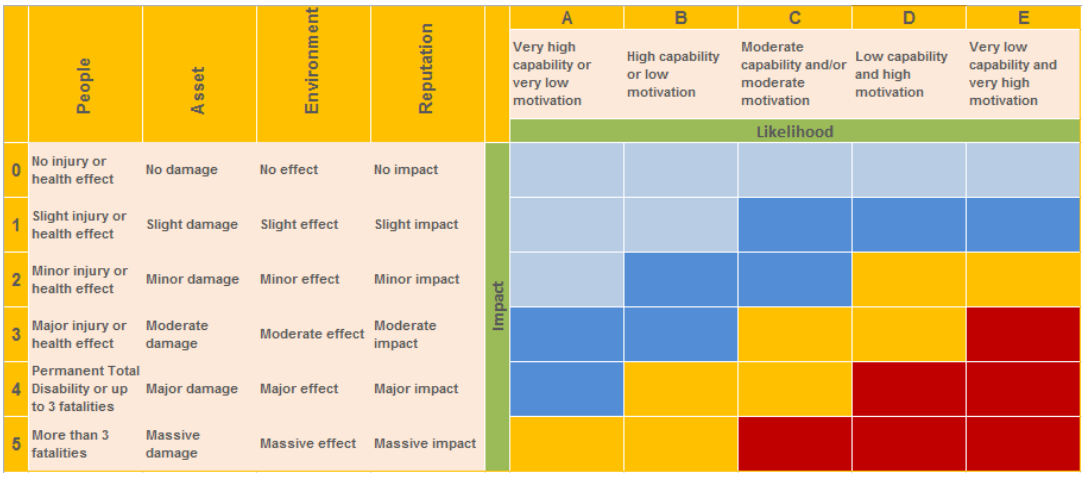

B. The simulation of Residual Risk after all Practice Areas are applied takes into account the fact that we cannot fully eliminate risk, and as more technical controls are applied, it becomes harder to decrease the likelihood of the top event occurring or the impact after the top event occurs. In other words, as you add technical controls, the resulting effectiveness of each added control decreases. As an example, consider the top event that a device gets a generic virus which is spread from another network layer. If Anti-Virus is added as a barrier, this Practice Area would easily decrease the likelihood of infection by the virus. If all Network Architecture controls are added instead, this would also easily decrease the likelihood of the top event occurring. However, if Network Architecture is added when Anti-Virus is already fully implemented, this will most likely not decrease the likelihood directly proportional to the impact of the barriers implemented individually.

C. The Bow-Tie is structured in such a way that certain Practice Areas are applicable for reducing the likelihood of a top event (such as Portable Media), and others are applicable to reduce the impact of a top event (such as Incident Response and Backup and Restore), as this is generally how these Practice Areas are applied. However, it needs to be taken into account that this is not a steadfast rule, and some Practice Areas may also have a mitigating effect on the other side of the Bow-Tie than that for which they are most commonly applied. As an example, consider Network Architecture. If an external attack requires network hopping in order to infect an ICS Component with a generic virus (top event), Network Architecture would have an impact on reducing the likelihood of the top event. However, if that virus is configured to send control data back through the network to an external source, Network Architecture would likely also

reduce the impact of the top event after it occurs. Therefore, although we consider Network Architecture as a "Respond" Practice Area in order to compare the Practice Area to the Bow-Tie Model, it needs to be taken into consideration that Network Architecture may also have an impact as a barrier to reduce the likelihood of the top event as well. Dependent on the threat scenario, this concept needs to be applied for each Practice Area.

D. Preparatory Practice Areas (Governance, Workforce Development and Management of Change) are not considered individually in terms of effectiveness as barriers, as these Preparatory controls are required in order to fully benefit from the implementation of technical controls. It is difficult to measure their direct impact alone on reducing the likelihood or impact of a top event, but it needs to be assumed that they are fully effective in order to evaluate the effectiveness of the other Practice Areas as barriers.

E. Initial Risk for non-adversarial threat scenarios is determined by both taking into account the past and using experience and understanding of the threat to evaluate what the frequency would have been if there had been no controls in place, as it is assumed that the past was not a true baseline without any controls implemented and this may have prevented other instances in the past. In other words, a true baseline for Initial Risk would use the RAM to determine the likelihood of the threat scenario based on the past prevalence when no security barriers were in place. However, it is not realistic to believe that there were no security barriers in place previously, so an extrapolation of past prevalence should be used to determine Initial Risk.

Resulting Industrial Control System Risk Model

The Cyber Risk Assessment Methodology described above can be summarized by the following Process and visualized in a resulting Risk Model, showing all Threat Scenarios, Initial Risk and Residual Risk.

Figure 6: Cyber Risk Assessment Process

Using this process, the Risk Model will have all threat scenario permutations applied to all exposure permutations with all possible Top Events considered for an Initial Risk of the worst case feasible scenario, and a Residual Risk based on the effectiveness of the controls within each Practice Area. For example, if there are 5 possible top events for a particular threat scenario, there will be 3 exposure permutations evaluated for each top event, resulting in 15 evaluated worst case feasible scenarios for one threat scenario. This example does indicate that while the model is simplifying infinite threat scenarios down to a limited number, each threat scenario is used to evaluate multiple different scenarios.

To determine Initial Risk, use the applicable RAM matrix (Figure 3 for a Non-Adversarial scenario and Figure 4 for an Adversarial Scenario) to quantify the Initial Risk in terms of People, Asset, Environment and Reputation (ex. 1A, 2B, 4D, 4D).

To determine Residual Risk for each threat scenario, exposure, and top event combination, all Practice Area controls are applied as barriers prior to the top event to evaluate their effectiveness on whether they would limit the likelihood of the top event occurring. Reduce this to “high effectiveness”, “low effectiveness” or “no effectiveness” for the sake of simplicity. The effective Practice Areas will reduce the likelihood of top event occurring. Then apply these same Practice Area controls as barriers assuming the top event has occurred, and determine their effectiveness at limiting the impact of the worst case feasible scenario. The effective Practice Areas will reduce the impact on the RAM after the top event occurs. To determine the

Reduction to Likelihood and Impact based on the effectiveness of the controls, Assumption B needs to be re-examined. Essentially, there are diminishing returns on the addition of new barriers, and this assumption will need to be translated into numbers to determine the Reduction to Likelihood and Reduction to Impact. In order to put this assumption into numbers, a graph shaped as follows will be required for both Likelihood and Impact.

Figure 7: Risk Reduction Calculations – Likelihood

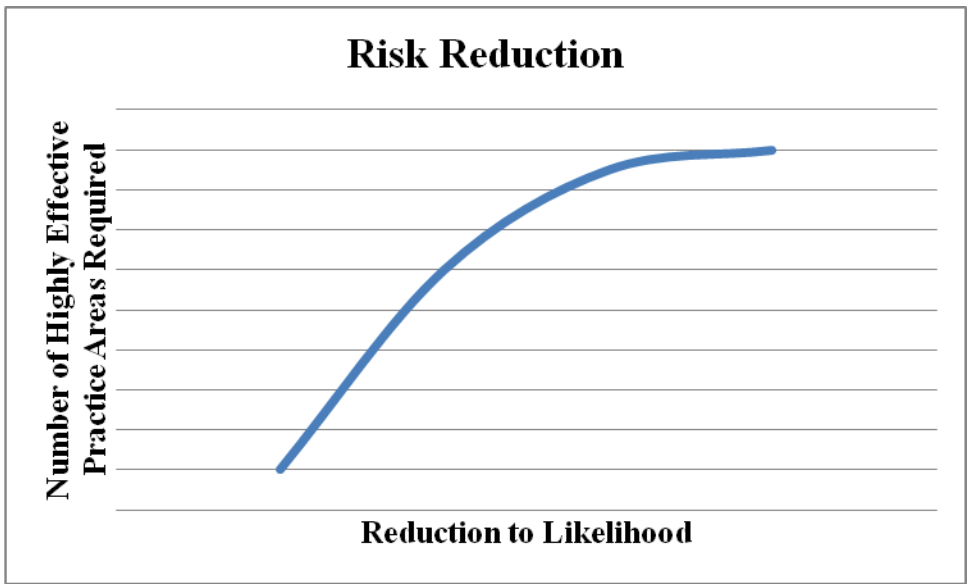

You can see that reduction to likelihood is lowered quickly when barriers are applied with high effectiveness, however the rate of increase slows as additional Practice Areas are added. The axis points will need to be determined by the organization and sanity checked using a real scenario or incident to determine whether the barriers are reducing the risk appropriately. Another consideration will be how a barrier with low effectiveness corresponds to a barrier with high effectiveness. Is a highly effective barrier twice as effective as a low effective barrier? Three times more effective? These considerations will be dependent on how conservative the resulting Risk Model needs to be.

The resulting likelihood and impact for each scenario will provide all Residual Risk for People, Asset, Environment and Reputation using the following calculations:

- Residual Likelihood = Initial Likelihood – Reduction to Likelihood
- Residual Impact = Initial Impact – Reduction to Impact

Ex. The Initial Risk for Reputation is 4D. The Practice Area barriers are determined to reduce the likelihood of the top event by 2 and reduce the impact by 1. Therefore, the Residual Risk for Reputation would be 3B.

For simplicity in the resulting model, you may choose to summarize each threat scenario by the highest Initial and Residual Risk from all of the permutations considered for that threat scenario.

Finally, to determine whether the resulting Standards are fit for purpose, the Risk Model can be used to justify entire Practice Areas by showing the change in the Residual Risk when they are removed from the model, or to justify elimination of an entire Practice Area if the change in Residual Risk is still considered ALARP. When evaluating individual controls within the Practice Area, consider the potential impact that removal or addition of that control would have on each ranking in the Risk Model, and determine whether the change in Residual Risk is worth the cost to implement or the cost savings to remove.

Furthermore, the Risk Model may also be used to evaluate a deviation request from the industry-wide ICS Security Standard, or to justify a control de-selection if the risks are different within a particular line of business, region or facility. Refer to Appendix B for an explanation of this process. This process allows for more educated decisions regarding the cost of control implementation in comparison to the determined Residual Risk.

5. Future Direction

An added complication to defending this Cyber Risk Assessment Process and resulting Risk Model is that the dynamic nature of the cyber threat landscape in conjunction with the unreliable incident data provided within the industry does not allow for a thorough comparison of the evaluated Initial and Residual Risks from the Risk Model with past incident frequencies and deemed impact. As this space evolves, and if reliable incident data was shared within the industry, the described Cyber Risk Assessment Process and Risk Model could be bolstered by this data. More importantly, the data could be used to justify the chosen axis numbers used to determine Risk Reduction in Figure 6.

While there continues to be reputational implications to sharing ICS cyber incident data within the industry, it is still an industry step in the right direction to have a shared methodology or process in assessing risk. It will make the industry stronger, as a whole, if the evaluation of the cyber threat landscape and resulting risk is conducted in a similar manner, ensuring that similarly strict ICS Security Standards are applied consistently around the world throughout the industry.

6. Results

The final result of the Cyber Risk Assessment Process is a standardized Risk Model for quantifying risk by evaluating all possible threat scenario and exposure combinations. The Risk Model is both easy to understand and explain because it is summarized into relatable threat stories and evaluated according to widely known Engineering concepts such as the Bow-Tie Model and Risk Assessment Matrix.

The Risk Model described above may not cover every requirement or diversity of condition within every location, region or industry; however, the described Cyber Risk Assessment Process may be used to justify the adaptation of controls to manage individual risks to ALARP. This Risk Model does not suggest that all risk is eliminated by implementing the required controls. Instead, it is a tool to justify the resulting ICS Security Standard, and estimate the Residual Risk based on that Standard. Due to the dynamic nature of ICS cyber threats, the Risk Model cannot be used to prove that a facility has removed all risk, as this is dependent on the current threat landscape, the operational implementation of the controls, and the design effectiveness maturity. This Risk Model is based on experience in the industry, and provides a best estimate of the Residual Risk to a location based on the assumption that all security controls are implemented with reasonable effectiveness. Note that implementation of controls as securely as possible may never be realized, and is an ongoing process that should be assessed with regular audits.

While the resulting Risk Model is not small, it successfully limits the threat scenarios down to specific permutations, and the method of using binary variables effectively addresses scrutiny regarding whether all possible threat scenarios have been considered. Each possible threat scenario has been addressed in the model, however the worst case feasible scenario may need to be adjusted as new threats are incurred which were not originally thought possible or feasible. Therefore, this Risk Model should be evaluated annually for its legitimacy based on the current threat landscape. The Residual Risk in the Risk Model will change as the threat

landscape changes, and the new Residual Risk will indicate whether new or more effective controls are required to manage the risks to ALARP.

Appendix A: Sample Risk Model Template

This sample template may be used to create a Risk Model based on the Cyber Risk Assessment Process described in this whitepaper.

Threat Scenario

- Example provided below for 1 of the 8 possible threat scenario permutations, and tied to a relatable story (Threat Story 1):
  - Agent (Ag.): Direct (Dir.)
  - Authorization (Auth.): Internal (Int.)
  - Motivation (Mot.): Adversarial (Adv.)

Exposure:

- All 3 permutations applied in example below to each applicable Top Event:
  - Business (Bus.): With Safety Critical Elements (With) or Without Safety Critical Elements (W/O)
  - System (Sys.): Essential (Ess.) or Safety System (Saf.)

Top Event:

- Example provided below has 2 applicable Top Events (TE1 and TE2)

Worst Case Feasible Scenario (assuming no barriers in place):

- Each combination has its own Worst Case Feasible Scenario (WCFS). Ex. Fail safe shutdown of Safety Systems.

Initial Risk:

- Each combination has its own assessed Initial Risk based on its WCFS. For example, WCFS1 would result in:
  - Initial Risk People for WCFS1 (P1)
  - Initial Risk Asset for WCFS1 (A1)
  - Initial Risk Environment for WCFS1 (E1)

    - Initial Risk Reputation for WCFS1 (R1)
- Threat Story 1 Initial Risk can be summarized by the worst Initial Risk from its applicable combinations:
    - Initial Risk People (PI)
    - Initial Risk Asset (AI)
    - Initial Risk Environment (EI)
    - Initial Risk Reputation (RI)

Control Effectiveness:

- Example provided below assumes there are only two Practice Areas in the ICS Security Standard (PA1 and PA2)
- Each Practice Area is evaluated for its effectiveness as a barrier on both the Likelihood and Impact of the Top Event:
    - No Effectiveness (N)
    - Low Effectiveness (L)
    - High Effectiveness (H)

Residual Risk:

- Each combination has its own calculated Residual Risk based on the evaluated Practice Area effectiveness:
    - Residual Risk People for combination 1 (P1)
    - Residual Risk Asset for combination 1 (A1)
    - Residual Risk Environment for combination 1 (E1)
    - Residual Risk Reputation for combination 1 (R1)
- Threat Story 1 Residual Risk can be summarized by the worst Residual Risk from its applicable combinations:
    - Residual Risk People (PR)
    - Residual Risk Asset (AR)
    - Residual Risk Environment (ER)
    - Residual Risk Reputation (RR)

Table 3: Sample Risk Model Template

Threat Scenario			Exposure		Top Event	Worst Case Feasible Scenario	Initial Risk				Control Effectiveness (Likelihood)		Control Effectiveness (Impact)		Residual Risk			
Ag.	**Auth.**	**Mot.**	**Bus.**	**Sys.**			**P**	**A**	**E**	**R**	**PA1**	**PA2**	**PA1**	**PA2**	**P**	**A**	**E**	**R**
Threat Story 1							PI	AI	EI	RI					PR	AR	ER	RR
Dir.	Int.	Adv.	With	Ess.	TE1	WCFS1	P1	A1	E1	R1	N/L/H	N/L/H	N/L/H	N/L/H	P1	A1	E1	R1
Dir.	Int.	Adv.	With	Saf.	TE1	WCFS2	P2	A2	E2	R2	N/L/H	N/L/H	N/L/H	N/L/H	P2	A2	E2	R2
Dir.	Int.	Adv.	W/O	Ess.	TE1	WCFS3	P3	A3	E3	R3	N/L/H	N/L/H	N/L/H	N/L/H	P3	A3	E3	R3
Dir.	Int.	Adv.	With	Ess.	TE2	WCFS4	P4	A4	E4	R4	N/L/H	N/L/H	N/L/H	N/L/H	P4	A4	E4	R4
Dir.	Int.	Adv.	With	Saf.	TE2	WCFS5	P5	A5	E5	R5	N/L/H	N/L/H	N/L/H	N/L/H	P5	A5	E5	R5
Dir.	Int.	Adv.	W/O	Ess.	TE2	WCFS6	P6	A6	E6	R6	N/L/H	N/L/H	N/L/H	N/L/H	P6	A6	E6	R6
Threat Story 2							PI	AI	EI	RI					PR	AR	ER	RR

Appendix B: Risk Model Options

The same Bow-Tie Method for conducting Cyber Risk Assessments for Industrial Control Systems can be used to create a Risk Model specific to a line of business, region or facility which has different operating conditions and requests a control de-selection or deviation from the ICS Security Standard. A control de-selection request would be appropriate if, for example, a specific facility does not have Process Safety risks or Safety Systems. In that case, the exposure level permutations involving Safety Systems can be removed from the Risk Model to determine whether all controls are required to maintain ALARP Residual Risk. A de-selection of controls may be granted from the industry-wide ICS Security Standard if the Residual Risk remains ALARP without implementation of the controls. In other words, the risk does not change based on not implementing those specific controls, and the facility maintains the same Residual Risk as all other facilities at a reduced cost of control implementation.

In contrast to a control de-selection, a deviation request could be evaluated by lowering Practice Area effectiveness in the Risk Model to determine how the Residual Risk is impacted by not implementing the controls within the requested deviation. A deviation would indicate that there is additional Residual Risk which is accepted, typically due to the justification of the cost to implement the controls or because there is a remediation plan and timeline already in place.

Appendix C: Definitions

Table 4: Definitions

Term	Definition
Threat Scenario	The combination of agent, authorization and motivation that may result in a top event with the premise that no control barriers are in place. Example: Employee of a supplier uses removable media to run software on an ICS device and unintentionally infects the device with a generic virus.
Actor	The individual or group carrying out the threat scenario. Example: vendor employee or Anonymous member
Agent	The medium through which your top event materializes. If the threat is materialized through someone else that is coerced to knowingly or unknowingly carry out a threat on another's behalf, the agent is still determined by how interaction with the target occurs.
Indirect	The top event materializes through the network. Example: Adversary logs on through the network
Direct	The top event materializes through removable media or physical access. Example: Adversary launches an attack physically at the computer by typing commands on the keyboard
Authorization	How access is gained to the system in order for the actor to use the agent to compromise process. This implies that the abuse of authorization is categorized as internal (i.e. rebooting a computer which you should not have been authorized to reboot) because the individual had the ability to carry this out using not publicly accessible access.
Internal	The actor has some form of access into the organization which is not publicly accessible. Examples: Supplier, vendor, contractor, or employee
External	The actor does not have access into the organization, or only has access to public information/locations. Examples: Hackers or Anonymous without internal access/information

Motivation	The intention behind the actions of the actor.
Adversarial	Individual, group, organization, or government that conducts or has the intent to conduct detrimental activities (malice). Example: nation state attack
Non-Adversarial	Accidental or unintentional activities. Note that environmental and structural events (such as hurricanes, power failure, etc.), are not included within our definition because they are not within the ICS Cyber Security scope of control. Example: human error, equipment failure
Top Event	The first cause, which has a direct or indirect potential consequence to the process. Example: An actor gains unauthorized access to a system – in and of itself this does not cause harm, however it has a potential adverse effect on the availability, integrity, or confidentiality of the process.
Exposure	The combination of Business and system type which are exploited by the threat scenario.
Business	Line of business exposed to the threat scenario, separated out by those business lines which could have Process Safety consequences and those which don't.
With Safety Critical Elements	Continuous operation facilities with higher Process Safety consequences due to the way in which the hydrocarbons are processed.
Without Safety Critical Elements	Businesses without Process Safety consequences if a threat scenario occurs without any security controls in place.
System Type	Device exposed to the threat scenario.
Essential	Systems which are required to sustain the operation of the facility. Examples: DCS systems
Safety	Systems required sustaining the health and safety of the process and its surroundings. Examples: SIS devices
Non-Essential	Systems which are not required to sustain the operation of the facility. Note that these are not included in the Risk Model because they may be used as a jump point to get to an Essential or Safety System; however these systems on their own will not result in a worst case feasible scenario. Examples: Domain controllers.

Consequence	The worst case feasible scenario if the top event materializes into a harmful event. This may have impact on the organization's reputation, assets, people or the environment if proper barriers were not place.
Worst Case Feasible Scenario Without Controls	The most feasible and highest level of impact to the organization (people, assets, environment, or reputation) based on the given threat scenario, exposure, and top event. Feasibility is evaluated periodically and is dependent on a number of factors such as actor, time, target. Likelihood is assessed on the assumption that there are no security controls in the model. This implies that you determine the likelihood based on an extrapolation of known experiences/events because facilities have always had some level of security controls in place.
Initial Risk	Evaluation of consequences to People (P), Asset (A), Environment (E), Reputation (R) based on the likelihood of the threat scenario (using the RAM). Note that this is used as a baseline, and may be extrapolated based on a scenario with no controls. It is assumed that the past does not truly reflect a zero control situation.
Control Effectiveness	The effect of the controls within the security practice area on the likelihood or consequence of the top event occurring.
Residual Risk	Evaluation of risk after implementation of controls.

Appendix D: Acronyms

Table 5: Acronyms

Term	Definition
ALARP	As Low As Reasonably Practicable
DCS	Distributed Control System
ICS	Industrial Control System
NIST	National Institute of Standards and Technology
PAER	People, Asset, Environment, Reputation
SIS	Safety Instrumented System

Appendix E: Authors

Rebekah Mohr – rebekah.mohr@gmail.com

Appendix F: Collaborators

The Cyber Risk Assessment Process and resulting Risk Model for Industrial Control Systems was designed in collaboration with Maarten Oosterink – Maarten.Oosterink@Shell.com

Appendix G: References

NIST Special Publication 800-30 Revision 1: *Guide for Conducting Risk Assessments*

Langer: Bound to Fail: *Why Cyber Security Risk Cannot Simply be "Managed" Away*

International Standards:

- ISO/IEC 27005:2011: Information Technology – Security techniques – Information security risk management
- NIST SP 800-30: Risk Management Guide for Information Technology Systems

Last Updated: September 19th, 2019

Upcoming SANS Training

Click here to view a list of all SANS Courses

SANS San Francisco Fall 2019	San Francisco, CAUS	Sep 23, 2019 - Sep 28, 2019	Live Event
SANS Dallas Fall 2019	Dallas, TXUS	Sep 23, 2019 - Sep 28, 2019	Live Event
SANS London September 2019	London, GB	Sep 23, 2019 - Sep 28, 2019	Live Event
SANS Kuwait September 2019	Salmiya, KW	Sep 28, 2019 - Oct 03, 2019	Live Event
SANS Tokyo Autumn 2019	Tokyo, JP	Sep 30, 2019 - Oct 12, 2019	Live Event
SANS Cardiff September 2019	Cardiff, GB	Sep 30, 2019 - Oct 05, 2019	Live Event
SANS Northern VA Fall- Reston 2019	Reston, VAUS	Sep 30, 2019 - Oct 05, 2019	Live Event
SANS DFIR Europe Summit & Training 2019 - Prague Edition	Prague, CZ	Sep 30, 2019 - Oct 06, 2019	Live Event
Threat Hunting & Incident Response Summit & Training 2019	New Orleans, LAUS	Sep 30, 2019 - Oct 07, 2019	Live Event
SANS Riyadh October 2019	Riyadh, SA	Oct 05, 2019 - Oct 10, 2019	Live Event
SANS Baltimore Fall 2019	Baltimore, MDUS	Oct 07, 2019 - Oct 12, 2019	Live Event
SANS October Singapore 2019	Singapore, SG	Oct 07, 2019 - Oct 26, 2019	Live Event
SANS Lisbon October 2019	Lisbon, PT	Oct 07, 2019 - Oct 12, 2019	Live Event
SANS San Diego 2019	San Diego, CAUS	Oct 07, 2019 - Oct 12, 2019	Live Event
SIEM Summit & Training 2019	Chicago, ILUS	Oct 07, 2019 - Oct 14, 2019	Live Event
SANS Doha October 2019	Doha, QA	Oct 12, 2019 - Oct 17, 2019	Live Event
SANS Seattle Fall 2019	Seattle, WAUS	Oct 14, 2019 - Oct 19, 2019	Live Event
SANS SEC504 Madrid October 2019 (in Spanish)	Madrid, ES	Oct 14, 2019 - Oct 19, 2019	Live Event
SANS Denver 2019	Denver, COUS	Oct 14, 2019 - Oct 19, 2019	Live Event
SANS London October 2019	London, GB	Oct 14, 2019 - Oct 19, 2019	Live Event
SANS Cairo October 2019	Cairo, EG	Oct 19, 2019 - Oct 24, 2019	Live Event
SANS Santa Monica 2019	Santa Monica, CAUS	Oct 21, 2019 - Oct 26, 2019	Live Event
Purple Team Summit & Training 2019	Las Colinas, TXUS	Oct 21, 2019 - Oct 28, 2019	Live Event
SANS Training at Wild West Hackin Fest	Deadwood, SDUS	Oct 22, 2019 - Oct 23, 2019	Live Event
SANS Orlando 2019	Orlando, FLUS	Oct 28, 2019 - Nov 02, 2019	Live Event
SANS Houston 2019	Houston, TXUS	Oct 28, 2019 - Nov 02, 2019	Live Event
SANS Amsterdam October 2019	Amsterdam, NL	Oct 28, 2019 - Nov 02, 2019	Live Event
SANS DFIRCON 2019	Coral Gables, FLUS	Nov 04, 2019 - Nov 09, 2019	Live Event
Cloud & DevOps Security Summit & Training 2019	Denver, COUS	Nov 04, 2019 - Nov 11, 2019	Live Event
SANS Paris November 2019	Paris, FR	Nov 04, 2019 - Nov 09, 2019	Live Event
SANS Sydney 2019	Sydney, AU	Nov 04, 2019 - Nov 23, 2019	Live Event
SANS Mumbai 2019	Mumbai, IN	Nov 04, 2019 - Nov 09, 2019	Live Event
SANS Bahrain September 2019	OnlineBH	Sep 21, 2019 - Sep 26, 2019	Live Event
SANS OnDemand	Books & MP3s OnlyUS	Anytime	Self Paced