Creating a Cyber Risk Intelligence Framework

Integrating Best Practices and Standards

By Jack Freund – ISSA Senior Member, Pittsburgh Chapter

This article introduces the concept of a Cyber Risk Intelligence Framework, based on several best practices and standards, which includes the review of several life cycles for managing incidents, threats, and risk analysis.

Abstract

Turning incident data into risk intelligence requires careful planning and coordination between multiple teams, processes, and frameworks. This article introduces the concept of a Cyber Risk Intelligence Framework, based on several best practices and standards. This includes the review of several life cycles for managing incidents, threats, and risk analysis, as well as a detailed look at which variables are necessary for translating data from one life cycle to another. The framework is presented with the relationship between the variables identified.

The modern threat landscape is very complex and dynamic. Organizations are frequently attacked by threat actors of various skill sets using a disparate set of tools that changes daily. The control posture that worked so well to defend an organization yesterday may no longer be capable of preventing those same attackers the next day, depending on what new weaknesses are uncovered in the interim. With such a challenging operating environment, one might be ready to give up trying to create a prioritized list of things to defend and instead declare that we should "assume breach" and defend everything.

Those practiced in risk analysis will likewise argue that there is benefit to the nuance, indeed to prioritizing the information assets that need protecting, as there are far too many things to protect and too little resources with which to protect them. Even in very large financial organizations that are told they have a virtually unlimited budget, they still have to compete with other organizations for people, and time unfortunately will always remain a limited resource.

With this as a backdrop, it's beneficial for organizations to leverage threat intelligence data that the cybersecurity operations teams have in bettering their risk management operations. Likewise, the threat intelligence teams can be better armed to protect the organization if they are using data gathered by the risk team. This marrying of threat and risk data is the foundation of the risk intelligence framework model presented here. This is a result of the combination of four distinct frameworks, each with a different purpose but for which there is enough overlap to be mutually beneficial.

The four frameworks that form the foundation of this integrated cyber risk intelligence model are :

1. The incident response life cycle of NIST 800-61 rev2[1]
2. The threat intelligence cycle, adapted from Krizan's 1999 book, *Intelligence Essentials for Everyone*[2]
3. RAND's 2007 publication, *Using Risk Analysis to Inform Intelligence Analysis*[3]
4. The OpenGroup's cyber risk quantification (CRQ) standards, (OpenFAIR or just FAIR)[4] [5]

1 NIST SP 800-61 Rev. 2, Computer Security Incident Handling Guide, https://csrc.nist.gov/publications/detail/sp/800-61/rev-2/final.

2 Krizan, L. (1999). *Intelligence Essentials for Everyone*. Joint Military Intelligence College. https://apps.dtic.mil/dtic/tr/fulltext/u2/a476726.pdf.

3 Willis, H. (2007). *Using Risk Analysis to Inform Intelligence Analysis*. RAND. https://www.rand.org/dam/rand/pubs/working_papers/2007/RAND_WR464.pdf.

4 OpenGroup FAIR Risk Taxonomy Standard. https://publications.opengroup.org/c13k.

5 OpenGroup FAIR Risk Analysis Standard. https://publications.opengroup.org/c13g.

The remainder of this article will discuss where the integration exists between these frameworks and which variables can be useful to share to and from each.

The term *risk intelligence* can have varying meanings in different context, but the intent here is to develop a general purpose cyber risk intelligence framework, and as such the definition provided by Tillman is helpful for understanding the scope:

> The organizational ability to think holistically about risk and uncertainty; speak a common risk language; effectively use forward-looking risk concepts and tools in making better decisions; alleviating threats; capitalizing on opportunities; and creating lasting value.[6]

Incident response life cycle

The NIST 800-61 incident response life cycle is shown in figure 1. It is comprised of the following phases: 1) preparation, 2) detection and analysis, 3) containment, eradication, and recovery, and 4) post-incident activity. Once at the end of this four-step cycle, the outputs are intended to be fed back into the beginning of the life cycle and into other life cycles as will be explained later.

Incident response is one of those processes that benefits from proper prior planning, and the first phase of this life cycle articulates that well. Here the organization would develop the resources necessary for an appropriate response. This can include acquiring the right technologies to aid in response, configuring computing assets to provide appropriate evidence in case of an incident, and setting up the right triggers to alert the staff when an incident is in process. Lastly, this process step includes hiring and training the right personnel to know the organizational resources, capabilities, and processes in the case of a confirmed incident.

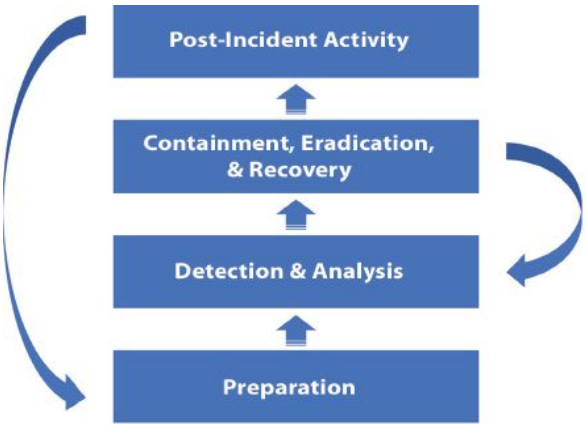

Figure 1 – Incident response life cycle

The second phase, *detection*, discusses the need for the organization to be able to know when an incident is unfolding in its environment. This means having the people and technologies in place (as per phase one) to be capable of reviewing precursor information and indicators to determine a priority for actioning in the next phase. This is due to the relatively high degree of false positives and negatives that may come out of monitoring and logging solutions. As a result, mature organizations establish a priority rubric for knowing when they need to escalate and respond quickly to certain indicators. This priority making is an excellent opportunity for integration with a risk intelligence model, as knowing the right risk-based priorities in this early phase of incident response is critical to managing risk well.

The third phase of this cycle brings about the closure of the incident. This means limiting the spread of the incident (quarantine), removing the attacker's access, payloads, and malware, and stopping any data exfiltration that is in process or soon to come. This also means attempting to recover systems and restore them to their state of operations prior to the incident. There should also be data gathering and evidence collection done in this phase, as it is sometimes necessary to hand this off to law enforcement or the human resources and legal departments for disposition and further actioning.

The final phase of the incident response life cycle, *post-incident activity*, really sets this up for integration with the others. In this, we are looking for a thoughtful collection of data and preparation of the resultant information for consumption by others. This can include a step-by-step recitation of what happened, as well as analysts' assertions about what this means to the organization. It can also include a forecast that can be used by others to know where in the organization to look for this type of incident in the future, as well as actionable next steps to help prevent this from happening again.

This step is really what turns an incident into actionable threat intelligence for consumption up and down the organizational chart. The level to which this goes in an organization may mean the removal of technical details but will always include a narrative section that summarizes what happened, turning it into a story that helps decision-makers understand what happened and adds credence to the analysts' assertions about what should be done next to prevent it. This information can be used as inputs back into phase one to aid in preparation for the next incident. This internally generated threat intelligence can also be used, along with similarly consumable threat intelligence from outside sources, as a critical input into the next life cycle for integration into the cyber risk intelligence framework.

Threat intelligence cycle

Threat intelligence is a system-two process as defined by Nobel prize-winner Danial Kahneman in his book, *Thinking, Fast and Slow*.[7] Briefly, in this book he details the kinds of decision making humans do and groups them into two categories: quickly made decisions that serve to protect us from harm and serve lower-level needs in Maslow's Hierarchy of Needs, and thoughtful decisions that come about as a result

6 Tilman, L. (2013). *Risk Intelligence: A Bedrock of Dynamism and Lasting Value Creation*. https://www.europeanfinancialreview.com/risk-intelligence-a-bedrock-of-dynamism-and-lasting-value-creation/.

7 Kahneman, D. (2013). *Thinking, Fast and Slow*. Farrar, Straus and Giroux.

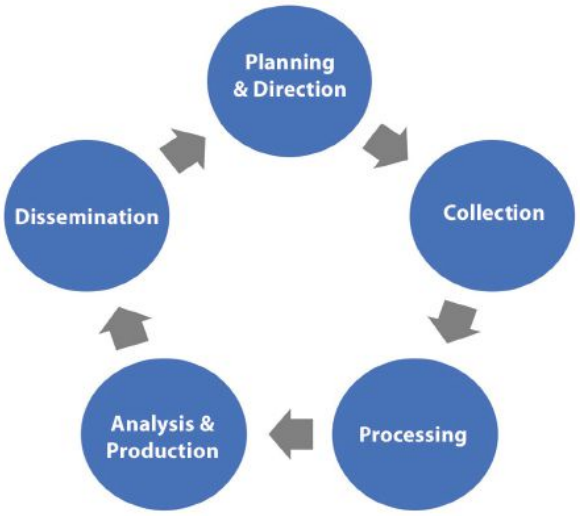

Figure 2 – Threat intelligence cycle

of spending time considering inputs and outputs, applying the right model for analysis, and curating the results. If this sounds time consuming, it is. System one thinking is so much easier and faster that we default to it often, which helps save us when we are imperiled, but can also give rise to poor choices at best, and war, bigotry, and racism at worst. Without question, threat intelligence needs to be a thoughtful endeavor that is designed to remove biases and carefully curate the information gathered to come to the most accurate and correct solution.

The general purpose threat intelligence cycle proposed by Krizan at the Joint Military Intelligence College (now called the National Intelligence University) gives us a cyclical life cycle designed to feed itself in an ongoing process of proposing a research question and them answering it. This life cycle can be seen in figure 2.

The first phase of this life cycle is to fully understand what questions the threat intelligence should answer and why. Further it requires of the analyst an assessment of which data needs collecting to answer this question. This phase should sound very familiar to any researcher, as it's the same thing that precedes any research project: what do we want to understand better and which data are necessary to do so. The overall risk intelligence framework will inform these research questions, so, for example, a threat intelligence researcher may want to better understand the types of attacks cyber criminals are using in her own organization and others like it.

Phase two is simply the collection of this data, or sometimes the identification of where this data may come from and putting in motion the technologies and processes necessary to generate and collect these data. Its at this point in phase three, once the necessary data are collected that the analyst will review the data, process it, and turn these raw data into actionable information. This means creating value-added reports in various mediums. Similar to the incident response life cycle, this includes the presentation of facts, conclusions, and forecasts about what can be expected next. Finally, this information is disseminated in various places, including the risk and threat analysts. This life cycle sits between the incident response cycle and the risk analysis cycle and as such has inputs into each.

Risk analysis life cycle

The basis for the life cycle presented here was taken from RAND's report on how to turn risk analysis into intelligence. For customization into the cyber-threat landscape, a few parts of the cycle have been modified (figure 3). Additionally, the necessary threat modeling variables have been taken from a fourth framework, the FAIR model for cyber risk quantification. This gives us the ability to manage threat and risk separately yet provide each with the actionable information it needs to operate smoothly.

In general, risk analysis is less sensitive to a threat agent's tactics and techniques than is threat intelligence. The reason for this is that while there may be various and sundry ways a threat agent may find a path into an organization's computing environment, the resultant risk tends to be more fixed. For instance, the ability of an attacker to leverage weaknesses in a retail organization's point-of-sale (POS) systems may change over time; however, the result will likely be the same (the compromise of and subsequent exfiltration of credit card data to conduct fraud). There are a few variables that are necessary for modeling risk well, especially quantitatively as we will see, and those are included here for review. Overall, its helpful to think of the risk analysis phase as an abstraction of the threat intelligence details that are ever so useful to the day-to-day blocking and tackling necessary to defend modern networks.

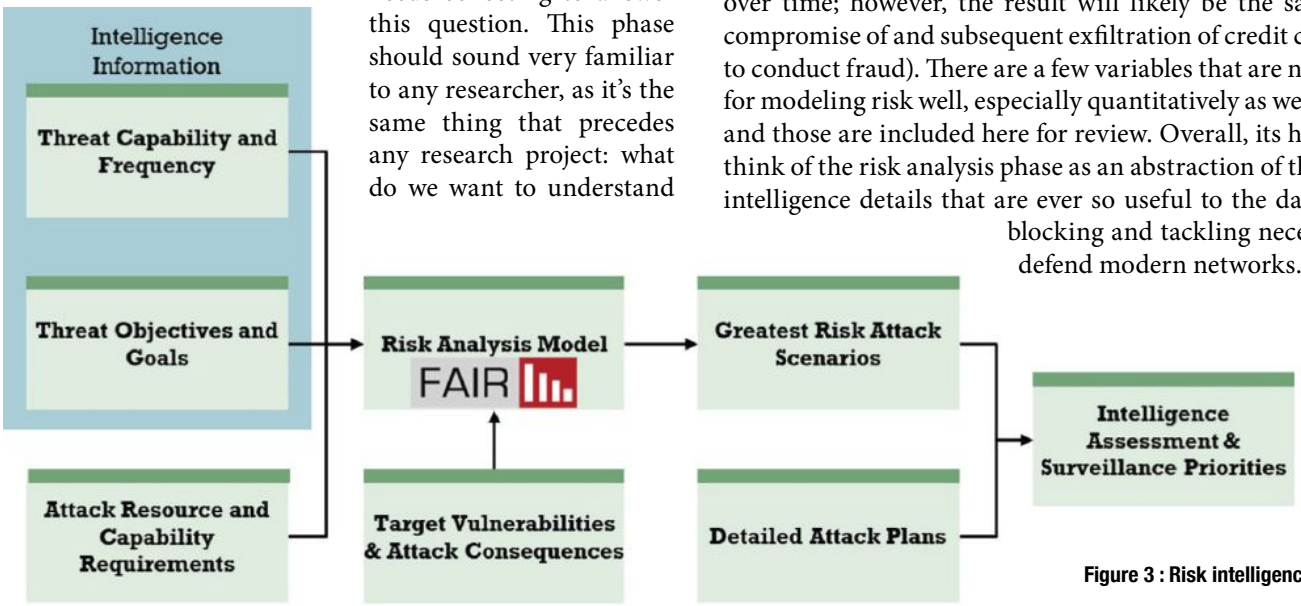

Figure 3 : Risk intelligence life cycle

The left side of the life cycle shows the inputs necessary for modeling risk, in this case the inclusion of relevant threat information. Leveraging FAIR, we can collect information necessary to model two variables called *Threat Event Frequency* (TEF) and *Threat Capability* (TCap). These two variables provide a model of how often threat agents are attacking (or if internal staff, then how often they make mistakes) and when they do, what force they can bring to bear. Accompanying this can be a threat-community profile that provides a common language for an organization to communicate about a threat community (figure 4).[8] This can detail information about the goals and objectives of the threat actors.

Data from the threat profile is funneled into these two variables, TEF and TCap, that are used as inputs into the FAIR model (figure 5). FAIR is designed to quantify cyber-risk loss exposure using economic measures, so the qualitative threat profile above will be translated into a measure of attempted loss frequency over time (TEF) and a measure of how capable the threat agents are (TCap). While frequency is a natively quantitative value, capability in the FAIR model is expressed in terms of the percentile level of the threat agent. For instance, attackers that can bring to bear the most strength in their attacks in terms of time, skills, and resources, can be said to be in the 99th percentile; thus they would have a TCap value of 99 percent. These values should also be stored along with the threat profile so that organizations can clearly communicate what these threat agents are to their organizations.

It's important to note that attribution is not a necessary prerequisite to building threat profiles for risk analysis purposes. Indeed, positive attribution is incredibly difficult and relatively unnecessary for the purposes of the risk intelligence process. Instead of strictly relying on an assessment that says an attack of a particular type has been positively correlated with a country's cyber capabilities, FAIR only asks whether

Motive	• Financial
Primary Intent	• Monetize proceeds of successful attacks
Sponsorship	• Usually none • Occasionally the beneficideary of state-sponsored intelligence (e.g., Russian KGB, Chinese government
General Targets Types	• Financial instruments and their issuing institutions
Specific Targets Types	• Financial services institutions • General public
Preferred Targets	• Liquid asset accounts • Payment cards
Concern for Collateral Damage	• Limited
Capability	• High degree of technological skill • Very high degree of social engineering skills
Personal Risk Tolerance	• Moderate to high

Figure 4 – Sample threat profile

8 Freund, J., Jones, J. (2014). *Measuring and Managing Information Risk: A FAIR Approach*. Portsmouth, NH: Butterworth-Heinemann.

ISSA International Web CONFERENCE

Click here for On-Demand Conferences

www.issa.org/?OnDemandWebConf

Legislative Aspects
Recorded Live: August 27, 2019

Paving the Way to a Passwordless Future
Recorded live: August 21, 2019

Beyond the Phish: Snapshot of End User Behavior
Recorded live: August 14, 2019

Privacy- GDPR a Year Later
Recorded Live: June 25, 2019

Passwordless Authentication
Recorded Live: une 12, 2019

Security-as-a-Service for Small and Medium-Sized Businesses
Recorded Live: June 5, 2019

Breach Response – Humans in Security
Recorded Live: May 28, 2019

What is a CASB and Why Do You Need It?
Recorded Live: May 22, 2019

Zero Trust – The Evolution of Perimeter Security
Recorded Live: May 15, 2019

Breach Report - Review the Various Breach Reports
Recorded Live: April 23, 2019

Practical Advice for the Proactive SOC: How to Escape The Vicious Cycle of React
Recorded Live: April 17, 2019

High Assurance Digital Identity in Zero Trust Architecture
Recorded Live: April 10, 2019

Threat Detection - Trends and Technology
Recorded Live: March 26, 2019

A WEALTH OF RESOURCES FOR THE INFORMATION SECURITY PROFESSIONAL

Figure 5 – FAIR ontology
[Graphic courtesy of the FAIR Institute]

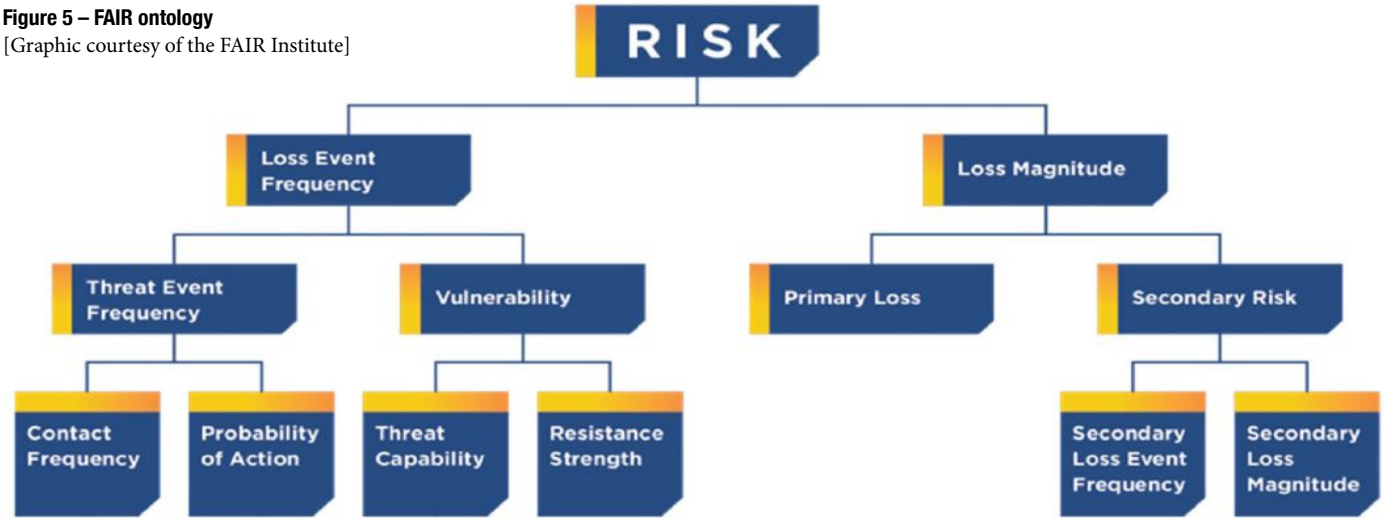

or not an attack of that type is likely to be a part of a broader, abstracted group (in this case nation-state attackers). Some users of FAIR may choose to model a specific country's threat capabilities, but it is not necessary for risk analysis.

FAIR also takes in information about the control state and economic impact to the organization to provide an accurate risk calculation. A complete treatment of the FAIR model is beyond the scope of this article; however, more information can be found at the nonprofit FAIR Institute's website, which promotes cyber risk quantification.[9]

Once several FAIR risk analyses have been completed, the organization will have a prioritized list of top risk scenarios. It's important to note that FAIR focuses on scenarios, which means there is a complete statement of loss, including the threat community, control weakness(es), and economic impact type that will be incurred (taken from the CIA triad). These top-loss scenarios can be used along with some forecasted attack plans (which can be aided by outputs from the threat intelligence life cycle) to focus the organization on the top ways attackers/insiders can bring about the realization of top-risk scenarios.

Cyber risk intelligence model

Each of the preceding life cycles has its strengths and usefulness in various parts of an overall cybersecurity program. Combining them together allows us to see what a mature cyber risk intelligence framework model can do for an organization. Figure 6 shows how these life cycles interact and where inputs and outputs can be mutually beneficial to each and the teams employing them.

9 FAIR Institute – https://www.fairinstitute.org.

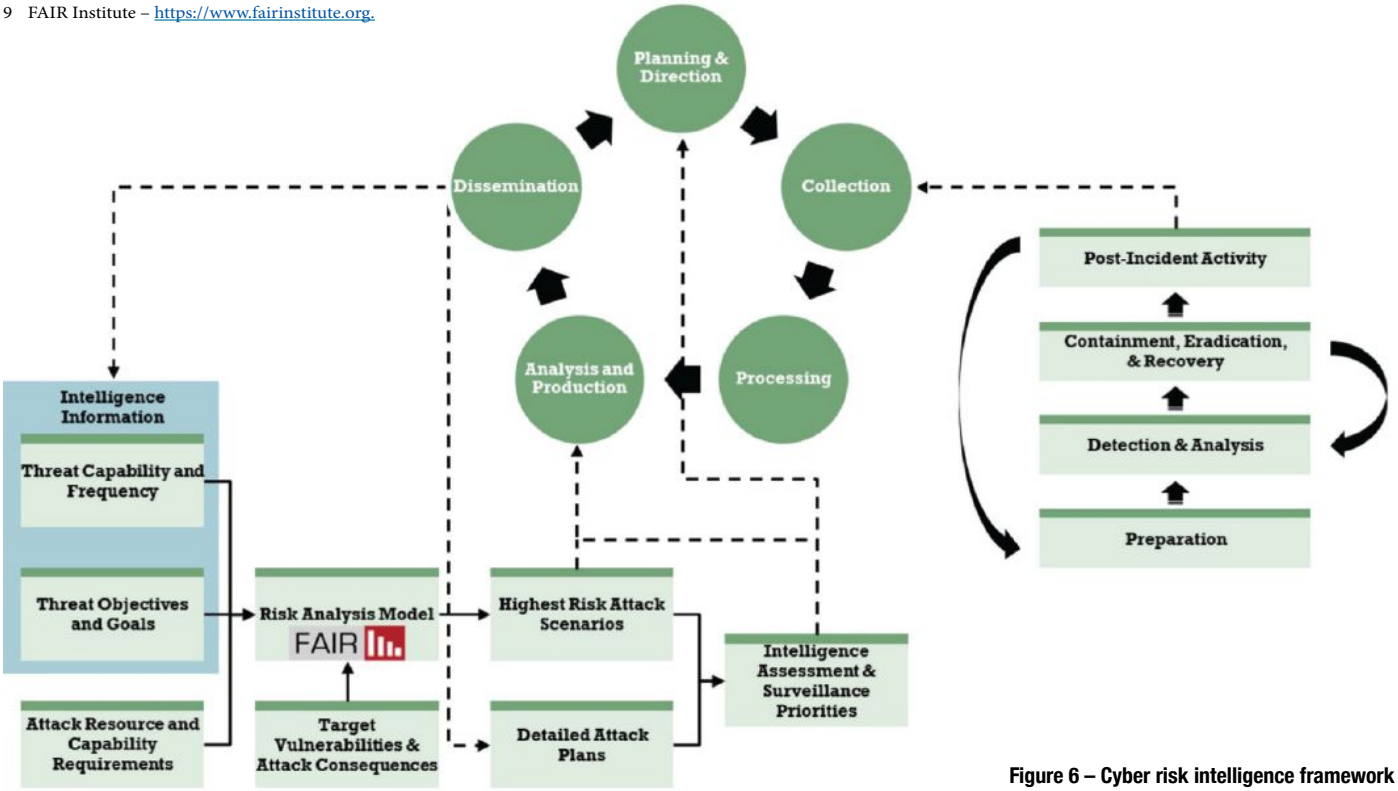

Figure 6 – Cyber risk intelligence framework

A summary of the interactions of these life cycles is presented here:

SOURCE	DESTINATION	COMMENT
Incident response life cycle: Post-incident activity	Threat intelligence life cycle: Collection	Internal and external incident reports can inform threat intelligence exercises and report writing
Threat intelligence life cycle: Dissemination	Risk analysis life cycle: Intelligence information	Threat intelligence reports can aid in the development of a threat profile and associated FAIR threat variables
Threat intelligence life cycle: Dissemination	Risk analysis life cycle: Detailed attack plans	Threat intelligence reports can be used along with top risk scenarios to better understand how to defend an organization against specific loss scenarios
Risk analysis life cycle: Highest risk attack scenarios	Threat intelligence life cycle: Analysis and production	Top-risk scenarios can be used as a part of the threat intelligence process of developing briefings and forecasting attacks
Risk analysis life cycle: Intelligence assessment and surveillance priorities	Threat intelligence life cycle: Analysis and production	The top-risk priorities, associated attack vectors, and control priorities can be used as a part of the threat intelligence process of developing briefings and forecasting attacks
Risk analysis life cycle: Intelligence assessment and surveillance priorities	Threat intelligence life cycle: Planning and direction	The top-risk priorities, associated attack vectors, and control priorities can be used as inputs into the threat intelligence research question creation

Conclusion

There is often a distance between the risk management function and the security operations and threat intelligence teams. This isn't on purpose; their work is just very different by nature. However, there is a need for these teams to collaborate in order to better understand each other's role and ground themselves in the reality that each deals with daily. The threat management team needs the business priorities that the risk team can offer, and the risk team needs to better understand threats to assess risk and business priorities. Using a cyber risk intelligence framework as the blueprint for collaboration gives explicit reporting responsibilities for each team. This can compel collaboration where there may not be any and create professional obligations that will build better teamwork and result in a higher quality of threat and risk intelligence products.

Many organizations are working towards improving their cybersecurity reporting to their senior management and boards. For these organizations, finding the right way to communicate the multitude of ways that bad things could happen as well as expressing potential cyber losses is a key concern. Coordinating narratives and sharing data between the cybersecurity operations teams and the cyber risk teams can help provide a more unified approach to not only those teams' daily work, but can bring unity in the end-to-end cyber risk stories being told to organizational leadership.

ISSA International Web CONFERENCE

ISSA Thought Leadership Series

Identities Are the New Security Perimeter in a Zero-Trust World

60-minute Live Event: Wednesday, September 18, 2019

10 a.m. US-Pacific/ 1 p.m. US-Eastern/ 6 p.m. London

In a recent Thales survey, two thirds of CISOs cited the increase in cloud service adoption, combined with a lack of strong security solutions, as the main reasons cloud services are the prime targets of attack. As organizations undergo digital and cloud transformation, CISOs and security officers are operating in a high-stress environment caused by security, compliance, and manageability challenges.

In this presentation we'll discuss how identities are becoming the new security perimeter in a zero-trust world and present best practices for implementing an access management framework that can help organizations remain secure—and scale—in distributed networking environments.

Generously supported by

THALES

CLICK HERE TO REGISTER.

For more information on these or other webinars:

ISSA.org => Events => Web Conferences

About the Author

Dr. Jack Freund, CISSP, CISA, CISM, CRISC, CIPP, PMP, is a leading voice in cyber risk measurement and management, using risk quantification to implement, mature, and sell information risk and security programs. Jack is currently serving as Director, Risk Science at RiskLens. The book Jack co-authored on quantifying risk, Measuring and Managing Information Risk: A FAIR Approach, *was inducted into the Cybersecurity Canon in 2016. Jack can be reached at jfreund@gmail.com.*

Copyright of ISSA Journal is the property of Information Systems Security Association, Inc. and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.