Effectively Leveraging Data in FAIR Analyses

Written for the FAIR Institute
Data Utilization Workgroup
October 2016

Jack Jones
EVP Research and Development, RiskLens

Effectively Leveraging Data in FAIR Analyses

With the advent of FAIR, organizations finally have a model that enables effective cyber risk measurement. The other necessary part of the equation, of course, is data. The good news on that score is three-fold:

- Calibration methods enable subject matter experts to provide good estimates
- Many organizations are awash in data from the security technologies that they use
- Many organizations have processes in place that can be useful sources of data

Of the three, the consensus is that, even with subject matter experts who are calibrated, it is preferable to limit reliance on subject matter estimates whenever possible due to the potential for cognitive bias and other human-related failings. Furthermore, the analysis process can be streamlined when good data exists and is appropriately normalized to the model. Besides, it simply makes sense to leverage security technology and process data when and where it exists. As a result, this document will provide guidance and examples to help organizations improve their FAIR-based risk analyses using these data sources.

NOTE: This document assumes the reader is familiar with FAIR-related concepts and terminology.

DATA-ASSISTED ANALYSIS VERSUS AUTOMATED ANALYSES

There is a continuum for using data in risk analysis. At the simplest level, data is leveraged to help inform just one of the factors in an analysis (e.g., TEF). At the other end of the continuum, all of the relevant factors in an analysis are provided by technology, which eliminates any need for human estimation, or for that matter, human intervention at all in the analysis. In this latter scenario, analyses can be fully automated, which would enable near-realtime risk analyses (or at least as realtime as the least current data element in the analysis).

There is, however, an important caveat to this last statement. Analyses can be automated once the scenarios being analyzed have been clearly defined, which includes defining all of the assumptions regarding how data will be applied within the analysis. This question of assumptions will be discussed in more detail further on in the document.

This document will focus on data-assisted analysis (the simpler end of the spectrum) because that is the most immediate need for organizations, and the most practical use-case. In doing so it will also begin to lay the groundwork for realizing automated analysis opportunities.

CHALLENGES WITH TECHNOLOGY DATA

For all the potential that data has for enriching risk analyses, there are challenges that have to be managed in order to use the data effectively. These challenges include:

- Analyst assumptions

- Raw versus “interpreted” data
- The relevance of historical data

ANALYST ASSUMPTIONS

It might seem logical to believe that having “hard data" for one or more of the risk factors in an analysis would reduce the need for an analyst to make assumptions. Unfortunately, in many cases this isn’t true. For example, let’s say that we want to use log data to support an analysis regarding malicious insider activity. Specifically, let’s say there is a key business application that processes sensitive customer information, and we need to estimate TEF. Although our logs may provide information regarding failed logins and rejected access attempts, the logs can’t tell us what percentage of those events were malicious in nature versus simple human error. They also often can’t tell us whether the actor, if malicious in intent, was looking to compromise confidentiality, integrity, or availability. The analyst has to apply some critical thinking skills, calibrated expert judgement, and, if they’re lucky, historical data regarding known malicious activity, to decide what percentage of those events should be considered threat events for this analysis.

The bottom line is that much of the data provided by technology is going to be ambiguous with regard to any specific risk scenario. In other words, an analyst still needs to figure out how or even whether to apply the data within an analysis. The assumptions underlying the analyst's decisions should also be identified and documented for independent review.

RAW VERSUS “INTERPRETED” DATA

Many security technologies have been designed to interpret the significance of the data they collect. For example, vulnerability scanners often use the CVSS model and/or a proprietary model to interpret the significance of any discovered deficiencies. Unfortunately, the results from those models usually cannot be applied within a FAIR analysis because those models define and calculate risk very differently than FAIR does. The good news is some of the *input* values to their calculations can be useful in a FAIR analysis. A good example are the Attack Vector metrics within CVSS (e.g., Local, Network, etc.). Even though the output from a CVSS analysis is typically useless from a risk measurement perspective, the CVSS Attack Vector metrics can be useful in estimating the relevance of a particular “vulnerability” to a specific risk scenario, which can help with TEF estimates.

It’s important to note too, that besides the fact that security technologies today calculate risk differently than FAIR, they also often don’t even fundamentally think of their input variables in the same manner. For example, CVSS uses its Attack Vector metrics in the “vulnerability” dimension of their analysis[1], whereas in a FAIR analysis, these metrics would usually be applied to help determine the TEF of a particular risk scenario (e.g. a network based attack versus a local attack, etc.). Consequently,

[1] Currently, CVSS does not explicitly include event frequency in any aspect of its risk formula

appropriately using metrics from these other models requires the analyst to think carefully about how they should be used in a FAIR context.

THE RELEVANCE OF HISTORICAL DATA

Data is implicitly historical in nature. As a result, an analyst has to consider the degree to which the past is likely to reflect the future. For example, to what degree does the TEF shown by the logs over the past year accurately reflect what the organization is likely to experience in the coming year?

Depending upon the scenario and the data point, you can often use historical data to drive one or more of the minimum, most likely, or maximum values for the specific data element (e.g., TEF, RS, etc.). When combined with other factors like strategic threat intelligence, pending new regulations, control improvements, etc., historical data can also be used to establish trends, which can influence the minimum, most likely, and maximum values of input variables.

Here again, critical thinking and calibrated professional judgment remains an important part of the analysis process.

USE-CASE EXAMPLES

This section will provide examples of how data can be leveraged to analyze several common risk scenarios. The intent is to illustrate the opportunities and challenges of leveraging data within not just these specific scenarios, but also in other scenarios that may be similar in nature.

These examples should not be considered comprehensive in their description of how the data can be leveraged. Readers are encouraged to extrapolate from these examples in order to gain optimum value from their data.

Anti-phishing solutions

One of the most significant security concerns for many organizations are phishing attacks. Consequently, in recent years technologies and service providers have sprung up to help organizations test, measure, and improve their ability to manage this concern. These technologies and service providers invariably provide data that describe the results of the testing, which can be helpful in performing FAIR analyses.

The very nature of this data speaks directly to the question of vulnerability. In other words, based on the testing, what percentage of attacks are likely to succeed? On the surface this seems like a very straightforward data point to use in an analysis — i.e., you'd simply take the percentage of successful attacks and plug that right into the vulnerability factor within a FAIR analysis. Unfortunately, depending on what you're trying to understand about phishing-related risk, it may not be that simple.

The data from these technologies/services often focus on two key data points: 1) the vulnerability of the average user within a population, and 2) which users fail repeatedly. For this discussion we'll focus on

when/how to use the first of these data points in two general phishing attack scenarios; 1) attacks targeted against specific personnel (spear-phishing), and 2) attacks against a broad population of personnel.

In spear-phishing analysis, you can use test results over time (e.g., the last 10 tests, etc.) for specific key personnel to inform a vulnerability estimate for attacks against those personnel. For example, let's say that a particular executive (Betty) is considered to be a likely spear-phishing target. Over the past year Betty has fallen victim to phishing tests 50% of the time. In it's simplest form, you could use 50% as the Most Likely value for vulnerability, and then some percentages above and below that for the Min and Max values, respectively, depending on how consistently Betty has performed. With that in mind, the data might also show that Betty seems to be improving because she was more susceptible earlier in the year than she has been lately. In that case, you might want to reflect her improvement as indicative of her future vulnerability with a lower set of vulnerability estimates (e.g., a minimum of 20% vulnerable, maximum of 50% vulnerable, and most likely of 30% vulnerable). Of course, you also would want to document your reasoning within the rationale portion of the analysis so someone reading the analysis would understand why the values used didn't align with Betty's average over the past year.

In an analysis of phishing against a broad population of personnel, the utility of phishing test data is likely to change, depending on how you structure the analysis. For example, you could structure the analysis such that the broad-based attack is considered a single threat event (albeit with a large attack surface). In this case, if your organization is like most, at least one person in the population will fall victim to the attack, which by definition means that the organization is 100% vulnerable to a broad-based attack. In this case, data regarding the percentage of personnel who fall victim can be best used to evaluate loss magnitude (i.e., the number of compromised systems that have to be cleaned up, the volume of information at risk, etc.). In this use-case, the value proposition for anti-phishing education/testing is reflected in a reduction in loss magnitude rather than loss event frequency.

The other way you could structure the analysis of a broad-based attack is to consider each phishing e-mail to each person as a distinct threat event. The phishing test results in this case can be used to drive vulnerability estimates in a more typical fashion — e.g., if 10% of personnel fell victim to the phishing test e-mails, it equates to 10% vulnerability. The challenge with this approach, particularly for large organizations subjected to frequent attacks, can be scale and complexity. Furthermore, the quality of your results often don't improve significantly for the extra analytic effort.

As you can see, appropriately using phishing test data depends strongly on what question/scenario you are trying to solve for, as well as how you structure the analysis.

NOTE: An additional consideration regarding phishing-related analyses is that personnel subjected to these attacks may be what we refer to as the point-of-attack but not the "objective of the attack." In such cases, these personnel are just the first step in an attempt for deeper penetration by an attacker. As a result, the loss event frequency of phishing attacks against personnel can be used to inform estimates for attacks against assets deeper within the environment (e.g., key databases). These are referred to as multi-level analyses, where a loss event at one level of abstraction in the environment becomes a potential threat

event for targets at a deeper level of abstraction. We'll see an example of this in the Malware section below.

Usage summary:

Vulnerability of individuals to phishing attacks
Vulnerability of populations to phishing attacks
Loss magnitude based on the potential number of compromised systems requiring forensics or other incident management efforts

Data leakage via e-mail

Another common concern for many organizations is the disclosure of sensitive or confidential information to unauthorized recipients via e-mail. In an attempt to manage this problem, many organizations have implemented Data Loss Prevention (DLP) solutions that identify, block, and/or report specific information being transmitted via e-mail. The data captured by these technologies can be used in various ways to support FAIR-based analyses.

For example, let's say that an organization wants to understand the risk associated with PII leaving the organization through e-mail. Fortunately, this organization has a DLP solution in place that captures whenever PII is contained within e-mail messages leaving the organization. On the surface it might be tempting to use DLP data regarding the frequency of these messages as the LEF value within an analysis — i.e., an assumption that each e-mail containing PII is considered a loss event. Before we start applying this data to any analysis though, it's a good idea to clearly understand the scenarios related to e-mail containing PII. For example:

- Malicious transmission to unauthorized recipients
- Accidental transmission to unauthorized recipients (e.g., the wrong recipient in the To: line of the message)
- Interception in transmission by unauthorized parties
- Non-malicious but intentional transmission to authorized recipients (e.g., as part of a business communication)

It is important to realize that the DLP data won't tell us which of these scenarios applies to any given e-mail containing PII. All it can tell us is that the transmission took place (or that an attempted transmission occurred, if it was blocked). The good news is that we may not need to differentiate. It may be perfectly acceptable to perform a single analysis that covers all four bases. Even so, it is important to have explicitly made and documented this choice, so that we can explain our reasoning if/when the question comes up from a stakeholder or other interested party.

Returning to the question of whether to treat each of these transmissions as a loss event, there are some important things to keep in mind:

- Within many organizations, e-mails transmissions containing PII occur every day
- The vast majority of these transmissions take place as part of normal (but perhaps ill-advised) business communications
- No loss materializes from the vast majority of these transmissions

As a result, to state that each of these transmissions (perhaps hundreds or even thousands per day in a large organization) is a loss event, is inaccurate. They may be violations of policy, and they may increase the potential for loss, but for the most part they don't materialize in loss.

One way to approach the analysis is to establish a current-state baseline of e-mail containing PII transmissions per time period (let's say quarterly). Establish another baseline value of known incidents where loss occurred involving PII in e-mail within the same time period. Derive from that an assumed percentage of incidents per volume of transmissions (e.g., .05% of transmissions resulted in loss). Using this as divisor, you can take any volume of transactions and derive an approximate loss event frequency. Clearly, although this isn't a perfect metric, it can be extremely useful as a defensible ballpark measurement of risk. It also enables a basis for comparisons as the volume of unauthorized transmissions changes over time (hopefully downward) due to better user education, better business processes, and other control improvements.

What I described above helps to resolve the frequency component of the risk equation. One of the nice things about DLP data is that it usually also provides insight into potential loss magnitude because it contains the volume and type of information within each transmission. The volume of sensitive records within any given transmission will in large part drive the potential loss magnitude should an incident take place. An organization can use a combination of its own loss history and loss data from organizations like RiskLens or the Verizon DBIR to help estimate losses for different volumes of records.

Usage summary:

Loss event frequency based on the number of unauthorized transmissions that result in loss
Loss magnitude based on the volume of sensitive information

Data discovery

Another potentially rich data source for risk analysis is the output from DLP solutions that scour an organization's systems looking for the existence of sensitive data. The data from this technology can be exceptionally useful when performing analyses that require estimates regarding the volume of sensitive information on various system types or storage media. For example, if an organization had performed a

scan of this nature against the workstations and laptops in their environment, they could readily make very good estimates regarding the minimum, most likely, and maximum number of records per device. Likewise, these data discovery technologies can identify network shares and other places where sensitive information might reside. Depending upon how granularly the scans were performed, the organization might even be able to differentiate records per device for different departments or geographies.

This type of information can be exceptionally helpful in making accurate loss magnitude estimates. For example, if an analysis is being performed to understand the risk associated with broad-based phishing attacks, understanding where and how much sensitive information resides on personnel laptops and desktops can strengthen loss magnitude estimates for these types of attacks. Furthermore, if an organization has significant variances in how much sensitive information resides on the systems within different departments, the organization can parse the analysis effort into separate analyses for greater precision and improved actionability. For example, if personnel within the Sales department all have some amount of sensitive customer data on their systems, while very few personnel outside of that department do, then it would make sense to do two separate analyses. This provides greater precision in analysis results, and also potentially provides the basis for requiring different controls on Sales systems.

Usage summary:

Loss magnitude based on the number of compromised systems requiring forensics or other intervention

Malware

(NOTE: In this section I use the term "anti-malware" as a generic reference that includes host-based as well as network-based anti-malware technologies. Also, in order to keep this section to a reasonable length, I limited the attack vectors under analysis to e-mail only. A full-fledged analysis might also want to include malware vectors such as malicious websites, USB media, guest computers, etc.)

Malware is a part of the risk landscape that most organizations have a lot of data for because most organizations (of any size) will have anti-malware technologies in place. In addition, many organizations will have anti-malware solutions at different layers of their technology architecture (e.g., perimeter systems, core servers, the network, personal systems, etc.), which increases the ways in which the data can be leveraged.

Let's say that an organization wants to understand how much risk it faces from e-mail borne malware. Given this objective, the organization will want to leverage data from its anti-malware technology to inform estimates regarding the frequency of malware attacks, the frequency of compromises from these attacks (which suggests the level of vulnerability the organization has) and, to some degree, the level of impact from these attacks. For example, the following data represents the known malware activity from week-to-week in this hypothetical organization:

Malware stopped at the e-mail server	
Week 1	1000
Week 2	950
Week 3	1113
Week 4	1022
Week 5	1013
Week 6	1054

Summary: Malware stopped at the e-mail server (per week)		
Minimum	**Most Likely**	**Maximum**
950	1013	1054

Malware detected on internal systems	
Week 1	15
Week 2	13
Week 3	21
Week 4	17
Week 5	31
Week 6	15

Summary: Malware detected on internal systems (per week)		
Minimum	**Most Likely**	**Maximum**
13	15	31

Malware infections requiring manual intervention	
Week 1	2
Week 2	3
Week 3	1
Week 4	2
Week 5	5
Week 6	2

Summary: Malware infections requiring manual intervention (per week)		
Minimum	**Most Likely**	**Maximum**
1	2	5

(NOTE: Within this example, we'll assume that the activity described above all occurred through the e-mail vector. Also, the numbers are for illustration purposes only and aren't intended to represent realistic malware activity for any particular organization.)

The simplest way to use this data is to focus on the last table — events where manual intervention was required to deal with a malware infection. With this approach, you could choose to ignore all of the other data and use the manual intervention summary values as the inputs for LEF in an analysis. This approach, however, fails to leverage some potentially important information.

Trending

Particularly within an active and dynamic problem space like malware, it is important to recognize trends over time. This is one reason why tracking malware data on a more frequent basis (like weekly) versus just annually, is beneficial.

Vulnerability

From the data in the tables above, we can measure the organization's vulnerability to malware at the perimeter.

Malware Vulnerability					
	Perimeter Data	Internal Detections	Total TEF	Loss Events	Vulnerability
Week 1	1000	15	1015	2	0.20%
Week 2	950	13	963	3	0.31%
Week 3	1113	21	1134	1	0.09%
Week 4	1022	17	1039	2	0.19%
Week 5	1013	31	1044	5	0.48%
Week 6	1054	15	1069	2	0.19%

Summary: Malware vulnerability (per week)		
Minimum	**Most Likely**	**Maximum**
0.09%	0.19%	0.48%

By taking the time to derive and monitor vulnerability an organization can identify and manage any important changes in vulnerability. For example, in week 5 the level of vulnerability was significantly higher than in other weeks. When this occurs, an organization can investigate why this was so. Was it a new virus strain that their anti-malware technology took a week to catch up to? Was it because of the introduction of a new set of systems that were installed without appropriate configurations? Was it due to an attack campaign by cyber criminals?

Regardless, without this type of data, it is more difficult to detect spikes and trends in the threat landscape and the organization's ability to defend itself. It also can be useful in setting risk-based KRIs.

Loss magnitude

Organizations can also use malware-related data to help inform the loss magnitude side of their risk analyses. Some of this information will come from tools, but much of it would come from the processes surrounding malware management. For example, the table below shows the losses involved in the manual interventions that were required for the various infections.

Manual Intervention Costs				
	Event	Person Hour Costs	Forensics Costs	Total Costs
Week 1	Event 1	$100	$0	$100
	Event 2	$100	$0	$100
Week 2	Event 1	$250	$0	$250
	Event 2	$200	$0	$200
	Event 3	$500	$5,500	$6000
Week 3	Event 1	$100	$0	$100
Week 4	Event 1	$150	$0	$150
	Event 2	$150	$0	$150
Week 5	Event 1	$350	$7,000	$7350
	Event 2	$100	$0	$100
	Event 3	$100	$0	$100
	Event 4	$250	$0	$250
	Event 5	$400	$2500	$2900
Week 6	Event 1	$200	$0	$200
	Event 2	$150	$0	$150

Summary: Manual Intervention costs (per event)		
Minimum	**Most Likely**	**Maximum**
$100	$100	$7,500

Additional columns could be added to capture other important pieces of information (e.g., volume of compromised records, whether the infections resulted in downstream attacks against other systems, the degree to which productivity was affected, etc.). These additional factors will help to flesh-out the loss magnitude side of the risk equation.

From a risk analysis perspective, this information helps risk analysts account for changes in the risk landscape over time, which should improve their ability to reflect probable levels of future exposure. Finally, this also enables an organization to pilot and measure the efficacy of, changes to their anti-malware solutions.

Usage summary:

LEF based on the number of manual interventions
TEF based on a combination of blocked infections and loss events
Vulnerability derived by LEF divided by TEF
Loss magnitude based on person hours and other costs resulting from manual intervention

Vulnerability scanning

Most companies today leverage vulnerability scanner technologies to help identify exploitable weaknesses in applications and systems. Many of these technologies employ the CVSS model to evaluate and report on the significance of any weaknesses that are discovered. The good news is that these technologies tend to be effective at identifying weaknesses. The bad news is that the CVSS model is not typically effective at accurately measuring the significance of the weaknesses it discovers. As a result, in many organizations the volume of "Critical" and "High Risk" CVSS findings is so high as to present a large signal-to-noise problem that limits the ability to prioritize effectively. There are, however, some useful data points (metrics) that come out of CVSS-based tools, that can be useful in FAIR-based risk analyses and for prioritizing weaknesses.

NOTE: This section will focus on a subset of the Base and Temporal metrics within CVSS and will not discuss every metric within the CVSS model. Also, some of the CVSS metrics not discussed here could, potentially, be applied when analyzing very specific scenarios. Generally speaking however, the metrics discussed here are more broadly useful when performing FAIR analyses.

There are three different value propositions that can be realized from a subset of the CVSS metrics:

- Support for estimating the vulnerability of a weakness
- Support for identifying which scenarios (and, thus, analyses) a weakness is relevant to
- Support for better TEF estimates

The discussion below will discuss how specific CVSS metrics can be used in each of these value propositions. Note that the significance of these metrics will vary from scenario to scenario. For this reason, the discussion below does not attempt to establish relevance levels for the CVSS metrics.

Analysts attempting to leverage CVSS data in their analyses will need to determine for themselves how much relevance to apply when using these metrics.

For complete information regarding the CVSS metrics and model, analysts are encouraged to read the documents that are available at the http://www.first.org/cvss URL.

Vulnerability

CVSS metrics that can be helpful when estimating vulnerability are relevant because they attempt to gauge the difficulty of exploitation faced by threat agents.

- The *Exploitability* metric attempts to measure the degree to which exploit techniques have been developed and made available within the threat landscape. From a vulnerability estimate perspective, weaknesses where exploit techniques are more broadly available and proven, should be easier for threat agents to leverage and thus could be considered more vulnerable — i.e., attacks require a lower level of attacker sophistication and capability.

- The *User Interaction* metric measures whether a user other than the threat agent needs to participate in order for an exploit to be successful. For example, many phishing attacks require that a user open a file or execute a command before the exploit can work. Weaknesses that require user interaction can be assumed to represent a lower level of vulnerability than those that don't.

- The *Attack Complexity* metric attempts to measure the inherent difficulty a threat agent faces in leveraging a weakness. Generally, this is due to circumstances surrounding a weakness and outside of the threat agent's control that increase exploitation complexity.

Threat Event Frequency

The vulnerability-related CVSS metrics discussed above also can be helpful when estimating TEF. These metrics tend to be relevant under one or both of two rationale: 1) the difficulty of exploitation reduces the population of threat agents that are capable of performing the attack, and 2) the level of difficulty reduces the number of threat agents who are willing to put forth the effort.

Scope

A couple of the CVSS metrics can be helpful when evaluating whether a weakness is relevant to any particular risk analysis scenario.

- The *Attack Vector* metric characterizes weaknesses in terms of where the threat agent has to be located relative to the target. Weaknesses that can be exploited over the network might be less relevant to analyses where the threat agent is local to the system. Conversely, weaknesses that can only be exploited by threat agents who are local to the system, may not be as relevant to analyses where the threat agents only have network access.

- The *Availability Impact* metric can be useful as a point of triage that includes or excludes weaknesses from an analysis. Specifically, if a risk analysis is focused on a confidentiality event, weaknesses with a high Availability value will not likely be relevant.

NOTE: CVSS also includes metrics for Confidentiality and Integrity. Unfortunately, they way CVSS defines and applies those metrics makes it difficult to apply them in risk analysis. These are examples of where, in specific analytic circumstances, CVSS metrics may be applicable but general use is problematic.

Usage summary:

Vulnerability estimates informed by various CVSS metrics
TEF estimates informed by various CVSS metrics
CVSS metrics can help determine the relevance of weaknesses within various risk analyses scenarios

Summary

At the end of the day, data from security technologies can be exceptionally useful in risk analysis. That said, before we can truly leverage the power of that data in broad and efficient ways, data has to be normalized on common foundational models. FAIR is the logical model for the risk component of that effort. Until then, risk analysts will continue to have their work cut out for them to effectively interpret and apply security technology data.

Even after security technology data has been normalized and integrated, there will always be limitations in how much of the cyber and technology risk landscape can be analyzed in a fully automated fashion, which means that subject matter expert judgment will remain a part of the analysis process. As a result, even organizations that are “data rich" will need to employ capable risk analysts in order to maintain a clear and accurate understanding of their risk landscape.