WHITE PAPER

A CLARIFICATION OF “RISKS”?

Written for the FAIR Institute
Cyber Risk Workgroup
January 2017

Jack Jones
EVP Research & Development, RiskLens

A Clarification of “Risks”?

People in the risk management profession routinely use the word “risk” in different ways. Although this may be fine in a non-professional setting, it presents significant challenges in terms of our ability to accurately and efficiently identify, measure, and communicate about risk.

In this short paper, I’ll provide some examples of common problematic usage, and I’ll explain why it’s problematic. I’ll then describe a simple set of criteria for determining whether or not something qualifies as “a risk.”

Getting on the same page

Let’s start out by making sure we’re on the same page about three things:

1. We care about identifying, measuring, and managing “risks” in order to help our organizations avoid adverse outcomes that can derail or impinge upon stakeholder objectives
2. In order to manage “risks” in a cost-effective manner, we need to be able to prioritize them effectively
3. Prioritization requires a normalized measurement — i.e., apples-to-apples

If you vehemently disagree with any of these, then the rationale behind what’s in the rest of this document probably won’t resonate with you. If you subscribe to the notion that risk can be both upside and downside, then for now just assume that I’m focusing on the down-side dimension.

One word – two contexts

“*How much risk does that risk represent?*”

In the sentence above, the word “risk” is used in two entirely different contexts. The first context, “*How much risk…*”, implies “risk” is essentially a volume measurement (it’s also an uncountable noun, for you language buffs).

While, the second context, “*…that risk…*”, implies “risk” is something you can refer to and say, “*That thing is a risk*” (a countable noun). And no, I’m not going to mention nouns again, but I’ve found the distinction to be useful[1].
The focus of this paper is to clarify the usage of “risk” in the second context.

[1] BTW — any time you put the word “a” in front of the word “risk”, you’re referring to risk in the second context. Likewise, you’re also using the second context if you make risk plural by adding an “s” to the end of it. If, however, you’re “measuring risk”, you are using the word in the first context.

Is that a risk?

When I give presentations I'll very often ask the audience to identify which of the following are "risks"?

- Disgruntled insiders
- Internet-facing web servers
- Untested recovery process
- Network shares containing sensitive consumer information
- Weak passwords
- Cyber criminals

With very few exceptions, audience members reply, "*All of them!*" But in fact, none of those bullet points represent "risks." Clearly, they all contribute to the risk landscape, but they represent three distinct parts of the risk landscape.

- Disgruntled insiders and cyber criminals are <u>threat communities</u>
- Untested recovery processes and weak passwords represent <u>potential control deficiencies</u>
- Internet-facing web servers and network shares containing sensitive information are <u>assets</u>

The minute we conflate these different elements of the risk landscape into an ambiguous bucket called "risks" we lose the opportunity to measure risk accurately, prioritize effectively, or communicate clearly. Here's why…

Measuring risk

Answering that earlier question, "*How much risk does that risk represent?*" clearly requires measurement, and the most common formula for risk measurement is to multiply (or somehow combine) Likelihood with Impact. With that in mind, let's substitute one of the bullet points above for "*that risk*" in our sentence.

"How much risk do Disgruntled Insiders represent?"

On the surface that seems like a perfectly fine question, and something we might want an answer to. Now, let's try to apply our risk measurement formula: Likelihood x Impact.

Likelihood of what? Impact of what? As it turns out, likelihood and impact both inherently apply to <u>an event</u>, and Disgruntled Insiders are not an event. They might perpetrate an event, but they are not an event in-and-of themselves, which means you can't apply a risk measurement formula to them directly. What you have to do first is define the event scenarios that involve them — something like: *A disgruntled insider leverages weak passwords to run off with sensitive consumer information residing on a network share*. Now that is a scenario you can measure the risk of using the measurement formula.

But did you notice the problem with that scenario? It included two of the other bullets in our list. In other words, in order to create a measurable scenario that involves the Disgruntled Insiders threat community, we had to include an asset and a control deficiency. In fact, I can create a number of different scenarios just using the elements in the earlier list of bullets, for example:

- Cyber criminals leveraging weak passwords to compromise an internet-facing web server.
- Disgruntled insiders affecting the availability of an internet-facing web server that has an untested recovery process.
- etc.

What this all adds up to is that, because the elements in our bullet list can be combined with one another in multiple ways (which creates overlaps and dependencies), prioritizing amongst them on a superficial basis is not realistic.

An example of where organizations commonly fall victim to "non-risk risks" is in risk registers. In every organization I've encountered over the past several years, their risk register has been made up of mostly (if not completely) control deficiencies, threat communities, and assets of concern — just like the earlier bullet list. And, as is the case in almost any risk register, each entry requires ordinal Likelihood and Impact ratings. But because these risk registry entries aren't event scenarios, the Likelihood and Impact values aren't tied to anything. Their basis is either completely ambiguous or based on fundamentally flawed analysis. When I've questioned the people who manage those risk registers about how they came up with their Likelihood and Impact values, their answers have always boiled down to something like one of the following:

- "Those are required fields, so we had to put something there."
- "Well, disgruntled insiders are pretty common (someone somewhere is always ticked off about something) so we rated it high likelihood, and they could act in a manner that has a catastrophic impact, so we rated it high impact."

In the first case, there is a better than decent chance that the assigned Likelihood and Impact values are not accurate. In the second case, the logic is badly flawed and the ratings don't stand up to even casual scrutiny. Regardless, these risk registers are not helping their organizations make well-informed decisions or manage risk effectively.

The exact same problem exists in every "Top 10 Cyber Risks" list I see organizations put together, which means those organizations almost certainly have not accurately identified the things that actually matter most, risk-wise.

It's all about loss events

At the end of the day, risk management is all about helping to ensure that our organizations experience a frequency and magnitude of loss that is tolerable to our stakeholders. And invariably, loss occurs from events — be they malicious, accidental, or natural. So what we need to identify, measure, communicate about, and manage are the frequency and magnitude of potential loss events. These are "risks." These are the things — the only things — we can apply Likelihood and Impact values too.

But what about non-compliance events? Aren’t the choice of a bad password, the misconfiguration of a firewall, or the failure to update policies “events” too? Yes, they are, so you can assign a Likelihood value to them. They aren’t loss events though — i.e., no asset has been harmed (yet) — so you can’t assign an Impact value to them. Any loss event that might occur as a result of a non-compliance event would have its own, perhaps very different, likelihood. You simply cannot automatically equate the likelihood of a non-compliance event with the likelihood of a downstream loss event.

What about something like reputation? Isn’t it commonly referred to as a “risk”? Yes, unfortunately, it is often mistakenly referred to as a “risk”. I wrote a blog post specifically on that issue some time ago (http://www.fairinstitute.org/blog/theres-no-such-thing-as-reputation-risk), but in a nutshell, the reason reputation damage isn’t a risk is because it’s an outcome of an event. In other words, some loss event occurred (a risk materialized) and one of the outcomes was damage to reputation.

Example risks

Having presented a case for what “risks” are and are not, in this section I’ll provide some examples of “typical” cyber and technology risks that organizations face. These examples should not be considered comprehensive, but rather represent a starting point for developing a list of risks your organization faces.

Risk (Loss Event)
Accidental disclosure of sensitive consumer information
Accidental disclosure of sensitive corporate information (e.g., a new market strategy)
Accidental disclosure of intellectual property
Data center outage
Data loss/destruction
Data integrity compromise
Online fraud
Malicious breach of sensitive consumer information
Malicious breach of sensitive corporate information
Malicious breach of intellectual property
Material financial misstatement due to IT deficiencies
Product/service degradation
Product/service outage
Product/service quality/integrity problem
Regulatory compliance failure

Obviously, these risks are very high level in nature. Most of the risk analysis you would perform would be at a more granular level, and this is where people often struggle. In fact, the learning curve in performing high-quality risk analysis tends to be with defining the loss event scenarios (the risks) that need to be analyzed. Which leads us to…

A taxonomy for defining risks

Being consistent in how you define risks is important for a couple of reasons. First, it helps you to avoid overlapping scenarios and double counting how much risk exists. Second, it helps you to avoid missing risks altogether.

With this in mind, I've found that a useful starting point is the age-old information security Confidentiality, Integrity, Availability triad. There are other, more complex extensions of this structure, but I haven't found them to be more effective for this purpose. The image below is a visual starting point for this taxonomy:

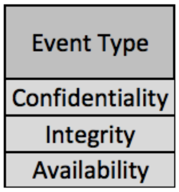

Event Type
Confidentiality
Integrity
Availability

From here we typically begin to parse-out the categories of assets at risk. This is where your organizational differences (industry, etc.) can begin to surface.

Another layer of abstraction (shown below) is typically necessary in order to achieve a good balance in terms of granularity. The "assets" at this layer of abstraction can be more detailed information types and/or business functions. Something to keep in mind here is that there is indeed a balance to be struck between enough granularity to be useful without straying into too much granularity. I like to refer to this as a "useful degree of precision". Not enough granularity means that you're not likely to gain meaningful insight into the landscape, and too much granularity begins to feel like counting grains of sand on the beach and can become overwhelming.

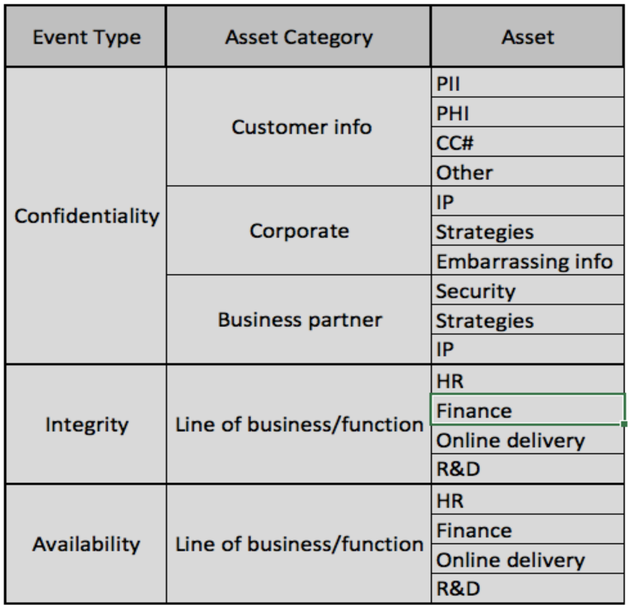

Event Type	Asset Category	Asset
Confidentiality	Customer info	PII
		PHI
		CC#
		Other
	Corporate	IP
		Strategies
		Embarrassing info
	Business partner	Security
		Strategies
		IP
Integrity	Line of business/function	HR
		Finance
		Online delivery
		R&D
Availability	Line of business/function	HR
		Finance
		Online delivery
		R&D

At this point, I typically like to begin parsing out the threat component of the landscape rather than getting more granular in parsing out assets. Your mileage, of course, may vary.

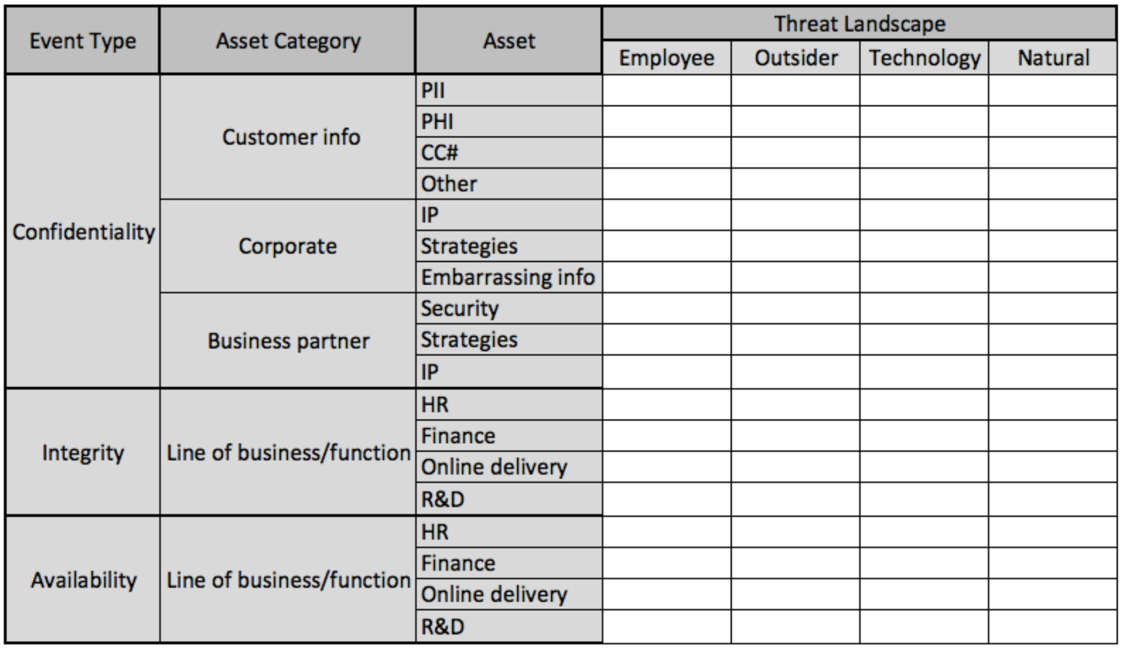

Event Type	Asset Category	Asset	Threat Landscape			
			Employee	Outsider	Technology	Natural
Confidentiality	Customer info	PII				
		PHI				
		CC#				
		Other				
	Corporate	IP				
		Strategies				
		Embarrassing info				
	Business partner	Security				
		Strategies				
		IP				
Integrity	Line of business/function	HR				
		Finance				
		Online delivery				
		R&D				
Availability	Line of business/function	HR				
		Finance				
		Online delivery				
		R&D				

At this point, you can continue to increase granularity in the threat landscape until you reach that balance point of granularity (by the way, "ET" under the Natural threat category represents "Extra-Terrestrial" events, like solar flares, etc.).

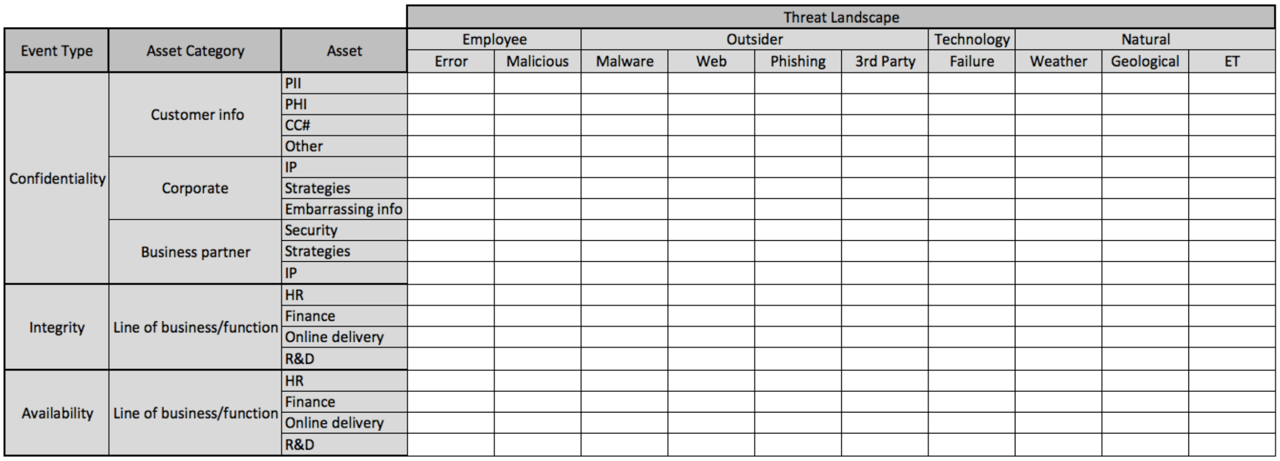

			Threat Landscape									
Event Type	Asset Category	Asset	Employee		Outsider				Technology	Natural		
			Error	Malicious	Malware	Web	Phishing	3rd Party	Failure	Weather	Geological	ET
Confidentiality	Customer info	PII										
		PHI										
		CC#										
		Other										
	Corporate	IP										
		Strategies										
		Embarrassing info										
	Business partner	Security										
		Strategies										
		IP										
Integrity	Line of business/function	HR										
		Finance										
		Online delivery										
		R&D										
Availability	Line of business/function	HR										
		Finance										
		Online delivery										
		R&D										

In the images below, I've broken out the threat landscape into one additional layer of abstraction (because the size of the matrix has gotten so large, I've broken the matrix up into three separate images). This is about as granular as I recommend going at first, as the matrix quickly explodes in size as you get more granular.

			Threat Landscape							
			Employee							
Event Type	Asset Category	Asset	Error				Malicious			
			E-mail	Web	Mail	Verbal	E-mail	Web	Mail	Verbal
Confidentiality	Customer info	PII								
		PHI								
		CC#								
		Other								
	Corporate	IP								
		Strategies								
		Embarrassing info								
	Business partner	Security								
		Strategies								
		IP								
Integrity	Line of business/function	HR								
		Finance								
		Online delivery								
		R&D								
Availability	Line of business/function	HR								
		Finance								
		Online delivery								
		R&D								

			Threat Landscape											
			Outsider											
Event Type	Asset Category	Asset	Malware				Web			Phishing			3rd Party
			Laptops	Desktops	Mobile	Servers	Apps	Servers	Network	Helpdesk	Executives	Administrator	
Confidentiality	Customer info	PII											
		PHI											
		CC#											
		Other											
	Corporate	IP											
		Strategies											
		Embarrassing info											
	Business partner	Security											
		Strategies											
		IP											
Integrity	Line of business/function	HR											
		Finance											
		Online delivery											
		R&D											
Availability	Line of business/function	HR											
		Finance											
		Online delivery											
		R&D											

			Threat Landscape						
			Technology	Natural					
Event Type	Asset Category	Asset	Failure	Weather			Geological	ET	
				Snow	Flooding	Wind	Earthquake	Asteroid	Solar flare
Confidentiality	Customer info	PII							
		PHI							
		CC#							
		Other							
	Corporate	IP							
		Strategies							
		Embarrassing info							
	Business partner	Security							
		Strategies							
		IP							
Integrity	Line of business/function	HR							
		Finance							
		Online delivery							
		R&D							
Availability	Line of business/function	HR							
		Finance							
		Online delivery							
		R&D							

There isn't anything sacred about how this taxonomy is defined, and you may very well come up with something more appropriate for your needs. It does, however, provide an example of how you can define your organization's risk landscape in a systematic manner, which can be extremely useful.

By the way, I wrote a multi-part blog post series on prioritization that leverages this taxonomy. You can find it on the FAIR Institute website (http://www.fairinstitute.org/blog/best-approach-to-prioritizing-risks-part-1).

Applying this to control deficiencies

All of this is well-and-good, but very often we're expected to evaluate how much risk a very specific concern represents — e.g., *How much risk does this audit finding represent*?

In order to answer a question like this, we have to identify the risks where the control in question is relevant. For example, deficient access privilege management may be relevant to several risks:

- Malicious insider acts against customer PII

  ✳confidentiality,

  ✳availability, or

  ✳integrity

- Accidental insider actions against customer PII

  ✳confidentiality,

  ✳availability, or

  ✳integrity

- etc…

Understanding how much risk the deficiency represents becomes a matter of evaluating and aggregating the risk associated with each of these scenarios. Getting this right requires two phases — 1) measuring the level of risk in the current deficient state, and 2) measuring the level of risk once the deficiency is remediated. The difference between the two represents how much risk the deficiency represents.

At first glance, this amount of analysis can seem daunting. There are ways of approaching it though, that are quite pragmatic, but that topic is beyond the scope of this paper and is something we can cover in the Cyber Risk Working Group. Regardless, unless approached in a systematic fashion such as this, the odds of measuring risk accurately and defensibly are significantly reduced.

Wrapping up

Cost-effectively managing risk requires that we measure it well and communicate about it clearly and consistently. This isn’t possible as long as we conflate different parts of the risk landscape. By defining “risks” as loss event scenarios, we significantly improve our ability to accurately measure and consistently communicate about risk.