

An Introduction to Factor Analysis of Information Risk (FAIR)

A framework for understanding, analyzing, and measuring information risk

DRAFT

Jack A. Jones, CISSP, CISM, CISA

LEGAL NOTICE

The contents of this white paper, and the FAIR framework are released under the:

Creative Commons Attribution-Noncommercial-Share Alike 2.5

More information on this license can be found here:

<http://creativecommons.org/licenses/by-nc-sa/2.5/>

For more information on FAIR, licensing, etc... please contact Risk Management Insight:

<http://www.riskmanagementinsight.com>

Or

Jack Jones

at jonesj1@riskmanagementinsight.com

Table Of Contents

Introduction	1
What's Covered...	2
Bald Tire Scenario	3
Scenario Analysis	4
The Bald Tire Metaphor	6
Risk and Risk Analysis	8
Risk defined	8
Purpose of risk modeling	9
Probabilities versus Possibilities	10
On being "wrong"	11
Summary	12
Risk Landscape Components	13
Threats	13
Assets	15
The Organization	16
The External Environment	16
Risk Factoring	17
Decomposing Risk	17
Loss Event Frequency (LEF)	18
Threat Event Frequency (TEF)	18
Contact	19
Action	20
Vulnerability	20
Threat Capability	21
Control Strength	21
Probable Loss Magnitude (PLM)	22
Primary Loss Factors	25
Secondary Loss Factors	27

Controls	31
Control dimensions	31
Control Lifecycle	33
Measuring Risk	34
Measurement Theory	34
Measuring Risk Factors	36
Measuring threat event frequency (TEF)	36
Measuring threat capability (Tcap)	37
Measuring control strength (CS)	38
Deriving Vulnerability	38
Deriving Loss Event Frequency (LEF)	39
Measuring probable loss magnitude (PLM)	39
Estimating loss	41
Deriving and Articulating Risk	44
Analyzing a Simple Scenario	46
The Scenario	46
The Analysis	47
Stage 1 - Identify Scenario Components	47
Stage 2 – Evaluate Loss Event Frequency (LEF)	48
Stage 3 – Evaluate Probable Loss Magnitude (PLM)	51
Stage 4 – Derive and Articulate Risk	55
Conclusions and Next Steps	57
Conclusions	57
Where to go from here	59
Appendix A: Basic Risk Assessment Guide	1
Stage 1 – Identify Scenario Components	2
Stage 2 – Evaluate Loss Event Frequency	3
Stage 3 – Evaluate Probable Loss Magnitude	8
Stage 4 – Derive and Articulate Risk	10
Appendix B: Glossary	1
Appendix C: Factoring Diagram	1
About the Author	1

Introduction

You can't effectively and consistently manage what you can't measure, and you can't measure what you haven't defined...

Ask a dozen information security professionals to define risk and you're certain to get several different answers. Pick up any information security book and you're likely to find that the author has used the terms risk, threat, and vulnerability interchangeably (they aren't the same thing). The simple fact is that our profession hasn't adopted a standard lexicon or taxonomy. The implications are not favorable, and many within the information security profession face the ramifications every day – for example:

- ! Marginalization in the organizations we serve
- ! Difficulty in convincing organizational leadership to take recommendations seriously
- ! Inefficient use of resources

As I see it, these issues practically scream “absence of credibility,” yet our most common response has been to complain, “they (the executives) just don't get it.” My recent observations suggest otherwise. Over the past several years it's become apparent to me that, far more often than not, executives DO get it. These are sharp people who live and breathe risk management as a significant and fundamental component of their jobs. It seems, instead, that the misalignment boils down to basic differences in definition and perspective. The executives are thinking “risk”, and we're

thinking “security” – two subtly, but critically different things, which will become apparent as we progress through this document.

The good news is that some within our profession have recognized the need to focus on risk, and have developed analysis processes and tools that take us in that direction. FRAP and OCTAVE® are a couple of the better-known examples. The unanswered challenge, however, is that without a solid understanding of what risk is, what the factors are that drive risk, and without a standard nomenclature, we can't be consistent or truly effective in using any method. FAIR seeks to provide this foundation, as well as a framework for performing risk analyses. It's important to note that much of the FAIR framework can be used to strengthen, rather than replace, existing risk analysis processes like those mentioned above.

Be forewarned that some of the explanations and approaches within the FAIR framework will challenge long held beliefs and practices within our profession. I know this because at various times during my research I've been forced to confront and reconcile differences between what I've believed and practiced for years, and answers that were resulting from research. Bottom line – FAIR represents a paradigm shift, and paradigm shifts are never easy.

Risk and risk analysis are large and complex subjects. Consequently, in writing this document I've had to balance the need to provide enough information so that risk concepts and the FAIR framework are clear and useful, and yet keep the length manageable. The result is what can best be described as an introduction and primer. For example, I've limited the scope to only include the human malicious threat landscape, leaving out threat events associated with error, failure, or acts of God. Some of the deeper, more complex elements of the framework also have been left out, and other elements have been brushed over lightly. Please accept my apologies in advance for the inevitable questions this introduction will leave unanswered. More thorough documentation is being developed. On the other hand, unanswered questions can be a good thing if they lead to dialog, debate, and additional research...

What's Covered...

The Bald Tire Scenario section will illustrate, through metaphor, some of the fundamental challenges facing the information security profession. It also briefly introduces some of the concepts that are fundamental to overcoming our challenges.

Before we can reasonably discuss the factors that drive risk, we first have to come to a common understanding of what risk is. **Risk and Risk Analysis** discusses risk concepts and some of the realities surrounding risk analysis and probabilities. This provides a common foundation for understanding and applying FAIR.

Risk Landscape Components briefly describes the four primary components that make up any risk scenario. These components have characteristics (factors) that, in combination with one another, drive risk.

Risk Factoring begins to decompose information risk into its fundamental parts. The resulting taxonomy describes how the risk factors combine to drive risk, and establishes a foundation for the rest of the FAIR framework. Note that we'll stay relatively high-level in our factoring to keep this from becoming a book.

The **Controls** section briefly introduces the three dimensions of a controls landscape.

Measuring Risk briefly discusses measurement concepts and challenges, and then provides a high-level discussion of risk factor measurements.

Bald Tire Scenario

As you proceed through each of the steps within the scenario below, ask yourself how much risk is associated with what's being described.

- ! Picture in your mind a bald car tire. Imagine that it's so bald you can hardly tell that it ever had tread. How much risk is there?
- ! Next, imagine that the bald tire is tied to a rope hanging from a tree branch. How much risk is there?
- ! Next, imagine that the rope is frayed about halfway through, just below where it's tied to the tree branch. How much risk is there?
- ! Finally, imagine that the tire swing is suspended over an 80-foot cliff – with sharp rocks below. How much risk is there?

Now, identify the following components within the scenario. What were the:

- ! Threats
- ! Vulnerabilities
- ! Risks

Scenario Analysis

Most people believe the risk is 'High' at the last stage of the Bald Tire scenario. The answer, however, is that there is very little probability of significant loss given the scenario exactly as described. Who cares if an empty, old bald tire falls to the rocks below?

Was my question about the amount of risk unfair? Perhaps, and I've heard the protests before... "*But what if someone climbs on the swing?*" and, "*The tire's purpose is to be swung on, so of course we assumed that somebody would eventually climb on it!*" Both are reasonable arguments. My point is that it's easy to make assumptions in risk analysis. In fact, some assumptions are unavoidable because it's impossible to know every conceivable factor within a risk scenario. However, assumptions about key aspects of the risk environment can seriously weaken the overall analysis.

The second point I'd like to make is that, from any group that goes through the Bald Tire scenario, I'll typically get several different descriptions of what constitutes the threat, vulnerability, and risk within the scenario. I've heard the frayed rope described as threat, vulnerability, and risk. I've also heard the cliff and rocks described as threat, vulnerability, and risk. The simple fact is that we, as a profession, have not adopted standard definitions for our terms. In informal discussions amongst ourselves, this may not always be a significant problem, as we typically understand what is meant by the context of the conversation. Consider, however, that physicists don't confuse terms like mass, weight, and velocity, and financial professionals don't confuse debit and credit – even in informal discussions – because to do so significantly increases the opportunity for confusion and misunderstanding. This is important to keep in mind when we're trying to communicate to those outside our profession – particularly to sharp executives who are very familiar with the fundamental concepts of risk – where misuse of terms and concepts can damage our credibility as professionals and reduce the effectiveness of our message.

A third point is that **you can't have significant risk without the potential for significant loss**. In other words, it doesn't matter how exposed to harm an asset is, *if the asset ain't worth much, the risk ain't high*. This is because **risk always includes a value component**. If it didn't, betting a million dollars would be equivalent to betting one dollar.

A final point is that there's a **tendency to equate vulnerability with risk**. We see a frayed rope (or a server that isn't properly configured) and automatically conclude that the risk is high. Is there a correlation between vulnerability and risk? Yes. Is the correlation linear? No, because vulnerability is only one component of risk. Threat event frequency and loss magnitude also are key parts of the risk equation.

So, what are the asset, threat, vulnerability, and risk components within the Bald Tire scenario? The definitions and rationale are described more specifically further on, but, simply stated:

- ! The asset is the bald tire
- ! The threat is the earth and the force of gravity that it applies to the tire and rope
- ! The potential vulnerability is the frayed rope (disregarding the potential for a rotten tree branch, etc.)

What about risk? Which part of the scenario represents risk? Well, the fact is, there isn't a single component within the scenario that we can point to and say, "Here is the risk." Risk is not a thing. We can't see it, touch it, or measure it directly. Similar to speed, which is derived from distance divided by time, risk is a derived value. It's derived from the combination of threat event frequency, vulnerability, and asset value and liability characteristics.

Having made an issue of terminology, the following paragraphs introduce and briefly discuss some basic definitions.

Threat

A reasonable definition for *Threat* is anything (e.g., object, substance, human, etc.) that is capable of acting against an asset in a manner that can result in harm. A tornado is a threat, as is a flood, as is a hacker. The key consideration is that threats apply the force (water, wind, exploit code, etc.) against an asset that can cause a loss event to occur.

Vulnerability

Although vulnerability is commonly recognized as a “weakness that may be exploited”, there’s more to it than that. We’ll dig into vulnerability more deeply later on, but for now we will leave it as a condition in which threat capability (force) is greater than the ability to resist that force.

You may have wondered why “potential” is emphasized when I identified the frayed rope as a potential vulnerability. The reason it’s only a potential vulnerability is that we first have to ask the question, “Vulnerable to what?” If our frayed rope still had a tensile strength of 2000 pounds per square inch, its vulnerability to the weight of a tire would, for all practical purposes, be virtually zero. If our scenario had included a squirrel gnawing on the frayed rope, then he also would be considered a threat, and the rope’s hardness would determine its vulnerability to that threat. A steel cable – even a frayed one – would not be particularly vulnerable to our furry friend. The point is that vulnerability is always dependent upon the type and level of force being applied.

Asset

In the context of information risk, we can define *Asset* as any data, device, or other component of the environment that supports information-related activities, which can be illicitly accessed, used, disclosed, altered, destroyed, and/or stolen, resulting in loss. The question is often asked whether corporate reputation is an asset. Clearly, reputation is an important asset to an organization, yet it doesn’t qualify as an information asset given our definition. Yes, reputation can be damaged, but that is a downstream outcome of an event rather than the primary asset within an event. For example, reputation damage can result from public disclosure of sensitive customer information, but the primary asset in such an event is the customer information.

Risk

The following definition applies regardless of whether you’re talking about investment risk, market risk, credit risk, information risk, or any of the other commonly referenced risk domains:

Risk – The probable frequency and probable magnitude of future loss

In other words – how frequently something bad is likely to happen, and how much loss is likely to result. As stated above, these probabilities are derived from the combination of threat, vulnerability, and asset characteristics.

Other Factors

So, where do the cliff and rocks fit into the risk equation? They aren't threat agents because they don't precipitate an event and, clearly, they aren't vulnerabilities that allow an event to occur. Consequently, these components can be considered *secondary loss factors* because their existence contributes to the magnitude of loss from an event. A real world example would be the fines and sanctions levied by regulatory agencies following an information security event. The regulations and regulators aren't the agents that commit a breach, so they aren't threats in the context of the event. They also aren't a technological, procedural, or other weakness that allowed the breach to occur. Nonetheless, they play a role in how much loss occurs and therefore must be included in our risk analysis. (Note, however, that there are scenarios in which regulators can be classified as threat agents – i.e., when they perform an audit.)

The Bald Tire Metaphor

Information risk management today is practiced as an art rather than a science. What's the difference? Science begins by analyzing the nature of the subject – forming a definition and determining the scope of the problem. Once this is accomplished, you can begin to form and then substantiate theories and hypotheses, which provide deeper understanding. This deeper understanding provides the means to explain and more effectively manage the subject.

Art, on the other hand, doesn't operate within a clearly defined framework or definition. Consequently, it's not possible to consistently explain or calculate based upon an artistic approach. A useful example is shamanism. The shaman rolls his bones or “confers with the gods.” He then prescribes a remedy based upon what his forefathers have passed down to him (“best practices”). Now, some shamans may be extremely intuitive and sensitive to the conditions within a scenario and may be able to select a reasonable solution on most occasions. But the shaman can't rationally explain his analysis, nor can he credibly explain why the cure works (or sometimes doesn't work). And, while we would like to believe that best practices are generally effective (as we tend to reuse what has been successful in the past), this may be a dangerous assumption. Best practices are often based on long-held shamanistic solutions, tend to be one-size-fits-all, may evolve more slowly than the conditions in which they're used, and can too often be used as a crutch – e.g., “I can't explain why, so I'll just point to the fact that everyone else is doing it this way.”

There is, however, no question that intuition and experience are essential components of how we do our jobs. The same is true for any profession. Yet these alone don't provide much traction in the face of critical examination, and are not strong formulas for consistency.

Putting Tread on the Tire

Recently, our profession has begun to pay a significant amount of attention to metrics. A word of caution – metrics and science are not the same thing. I can measure some parameter or count the instances of some event, but if I haven't developed a logical and rational understanding of the broader context within which the metric applies, all I have is numbers. Furthermore, in the absence of a fundamental understanding of the subject, it's far too easy to misinterpret and misuse the data. In order for metrics to be truly useful, we have to understand our subject.

We can't consistently and effectively manage what we can't measure – and we can't measure what we haven't defined. The first thing we need to do to shift from art to science is define our subject. What exactly is information risk? What are

the factors that make it up, and how do they relate to one another? After we've defined our subject, how do we measure it? How do we model and evaluate the complex risk scenarios we face? Finally, if we've managed to accomplish all of these things, how do we articulate risk to the decision-makers who need this information?

FAIR provides a reasoned and logical framework for answering these questions:

- ! **A taxonomy** of the factors that make up information risk. This taxonomy provides a foundational understanding of information risk, without which we couldn't reasonably do the rest. It also provides a set of standard definitions for our terms.
- ! **A method for measuring** the factors that drive information risk, including threat event frequency, vulnerability, and loss.
- ! **A computational engine** that derives risk by mathematically simulating the relationships between the measured factors.
- ! **A simulation model** that allows us to apply the taxonomy, measurement method, and computational engine to build and analyze risk scenarios of virtually any size or complexity.

As mentioned in the introduction, this document will brush over some components of the framework, and won't discuss others at all. For example, the measurement approach described in this document has been significantly simplified, the computational engine is not discussed in detail, and the simulation model isn't discussed at all. Nonetheless, the material provides a foundation for better risk analysis, and sets the stage for discussion, debate, and for subsequent documentation and training.

Risk and Risk Analysis

Our first challenge is to define the nature of the problem we're trying to solve – i.e., what is risk? This section will briefly cover the nature of risk and some simple truths about risk analysis.

Risk defined

Risk – The probable frequency and probable magnitude of future loss

There are three important things to recognize from this definition. First and most obvious – risk is a probability issue. We'll cover this in more detail throughout the document, so I won't belabor it now. Second – risk has both a frequency and a magnitude component. And third – the point I'd like to focus on here – is that this definition for risk applies equally well regardless of whether we're talking about investment, market, credit, legal, insurance, or any of the other risk domains (including information risk) that are commonly dealt with in business, government, and life. In other words, the fundamental nature of risk is universal, regardless of context. The good news is that risk concepts have been studied for generations within other professions, so a lot of good information is available. The not so good news is that we have, far more often than not, approached information risk as if it were somehow different from the other risk domains. This is one of the first hurdles we have to overcome if we hope to really understand our problem space.

Purpose of risk modeling

The purpose of any risk analysis is to provide the decision-maker with the best possible information about loss probabilities. Consequently, it's crucial that decision-makers accept the risk analysis methodology being used, and that the information resulting from the analysis is in a form that's useful to them. In this regard, the limitations of our traditional information security "risk analysis" methods will become clear as we progress through this document.

Limitations of risk analysis

Risk analysis is never perfect. The fact is, all risk analysis models are approximations of reality because reality is far too complex to ever model exactly. Nonetheless, by decomposing a complex subject into clearer, more readily analyzed components, we can understand and make reasoned judgments about the subject.

Any analysis model will be limited by the complexity of the subject, how deeply we understand the subject, and how well the model incorporates that understanding.

Evaluating risk analysis methods

FAIR is just one way of skinning the risk analysis cat. There are many dedicated people working to develop other methods with the same goals in mind. This is terrific news to those of us who are trying to be as effective as possible in managing risk. However, regardless of the methods you consider using, I encourage you to evaluate any risk analysis method on at least three points.

- ! Is it useful?
- ! Is it logical?
- ! Does it track with reality?

A methodology is useful when it meets the needs of the people making the risk decisions. If our analysis provides information regarding the effectiveness of controls in an environment, but decision-makers are looking for information regarding the probability of incidents and losses, then the methodology isn't useful. Likewise, if we provide "key risk indicator" metrics, but there is no clear linkage between the conditions described by the metrics and the probability of loss, then the metrics become little more than window dressing.

There are "risk analysis" methods in use today whose logic falls apart under close examination. Often, these methods call themselves risk analysis when what they're really analyzing is a subcomponent of risk (e.g., vulnerability or controls). Keep in mind, however, that these methods may be excellent at what they actually do, so don't disregard their value. It's simply a matter of recognizing what they do and don't provide. A simple way of identifying a bona fide risk analysis method is to determine whether it includes an analysis of threat frequency, vulnerability, and loss magnitude, and whether it treats the problem probabilistically. If one or more of these components is missing, or if the problem isn't treated as a probability, then logically, it isn't a risk analysis method.

The last point of consideration – tracking with reality – is especially critical. It's also a point that may be particularly challenging for our profession to come to terms with. Let's use, as an example, the controls condition of a typical corporate internal network. My experience has been that most corporate internal networks and systems are not very

well protected (at least relative to theoretical “best practices”), yet relatively few organizations have actually experienced a severe loss associated with this fact. Significant loss events do occur, no question of that, but they’re astonishingly infrequent given the prevalence of bad security practices. A realistic risk analysis should reflect this fact. In other words, an effective risk analysis of an average internal corporate information technology environment should, in most instances, identify very few truly high risk issues in spite of the fact that controls may not meet best practice standards.

Our profession has become very good at evaluating the technical controls in an environment – password length/complexity/history, whether appropriate access privileges are in place, how many layers of firewalls there are, how up-to-date the antivirus product is, whether intrusion detection exists, etc. We also know what the key non-technical controls are – whether policies exist, the level of user awareness, executive sponsorship of the security program, etc. But **controls are only part of the risk equation.** In fact, controls are only part of the vulnerability equation. We’ll cover the risk “equation” as we go along but, as I described in the introduction, vulnerability is always relative to the type and level of force being applied. If we only measure controls, but call the results “vulnerability,” then we’ve made significant assumptions about the level and type of threats involved. In order to effectively evaluate risk, all of the components of risk must be included.

When you’re evaluating any risk analysis method – FAIR included – ask hard questions. Make certain that it meets the needs of the decision-makers, that it’s logical, and that it aligns with reality.

Probabilities versus Possibilities

Possibility is a binary condition – either something is possible, or it’s not – 100% or 0%. **Probability** reflects the continuum between absolute certainty and impossibility.

Too often in my career, I’ve encountered executives and others who view the information security profession as being paranoid and full of “Chicken Littles proclaiming that the sky is falling.” Unfortunately, this perspective is generally well founded. We’ve tended to speak in terms of “it could happen” (possibility) rather than in terms that describe the probability of something happening.

The simple fact is that **risk is always a probability issue.** Consider the difference between playing Russian roulette with a standard six-cylinder revolver versus a semi-automatic. The possibilities are equal with either handgun – i.e., it’s 100% possible in both cases that the player would suffer a “negative outcome.” The probabilities, however, are significantly different. In the first case, assuming the revolver is loaded with a single bullet, the probability of a negative outcome is about 17%. In the second case, assuming a single bullet is loaded and chambered in the semi-automatic, the probability of a negative outcome is about 100% (it might, of course, misfire). Clearly, I’d rather not play the game at all, but if I had to choose between the two weapons, I’d much rather base my choice on an understanding of the probabilities, as opposed to just the possibilities. Decision-makers want and need the benefit of this same quality of information.

The natural concern, of course, is how we’re supposed to determine probabilities when there’s so little empirical data regarding information risk. I’ll go farther than that – not only is there very little information risk data, most of the data we do have isn’t credible. Here’s why. In order to establish credible conclusions from data, the data has to be reasonably accurate, current, and statistically significant as a sample. In the information risk realm, the accuracy of existing data has to be seriously questioned because it can’t be normalized against a standard taxonomy (i.e., because of our terminology challenges, one person’s “vulnerability” is another person’s “threat”, etc.). This absence of a taxonomic framework also

presents a significant challenge due to the sheer complexity and variety of our risk landscapes and the fact that the data we have today doesn't include details regarding the contributing factors for risk. For example, the annual CSI/FBI survey doesn't describe what specific conditions existed within the companies that did or didn't experience loss due to hacking. Consequently, we can't know whether common contributing factors existed in those companies that experienced loss, versus those that didn't experience loss. Furthering our challenge, the components within our risk landscape change so rapidly that the useful lifetime of data can be a problem.

An absence of good data doesn't relieve us of the obligation to deal with information risk as a probability issue. It's worthwhile to point out that non-data driven analyses have been successfully adopted for various applications, including medical diagnosis, missile targeting, rocket engine control, and marketing and investment, to name a few.

On being “wrong”

“Prediction is very difficult, especially about the future.”

(Nobel Laureate and nuclear physicist, Niels Bohr)

Many people become very uncomfortable with the notion of estimating probabilities, especially if they believe that they could be perceived as, and held accountable for, being wrong if the future unfolds differently than it seems they predicted.

Risk analysis is fundamentally all about establishing probabilities, and I can guarantee that if you do it often enough, at some point the future will unfold in a way that leaves open the possibility for others to perceive that you were wrong. Now, there always is the possibility that you were mistaken with your analysis – we're human, after all. But even if your analysis was perfect, the future is uncertain. A great example is rolling a pair of six-sided dice. Assuming the dice and roll aren't fixed in some manner, we can establish with a high degree of confidence that there's about a 2.7% probability of rolling snake-eyes, or, in other words, about once every thirty-six rolls. Now, if snake-eyes comes up on the first roll, does that mean that our probabilities were wrong? No! There's no rational reason to expect the dice to wait until the thirty-sixth roll to turn up snake-eyes. The key thing to keep in mind is that **establishing probabilities is not the same thing as foretelling the future.**

Decisions entail uncertainty

Any time we're forced to make a decision, we do so without perfect knowledge of what the outcome will be. We may be nearly certain, but nothing provides absolute certainty of what the future holds. This uncertainty introduces risk – i.e., the outcome of any decision may be undesirable.

The good news is that most businessmen and leaders are very aware that no guarantees exist and that the nature of business entails uncertainty and risk. They also understand the notion of probabilities and appreciate an analysis that articulates risk in those terms rather than in terms of “it could happen.” The “it could happen” analysis is generally looked upon as a “Chicken Little” position and is of little value to the person making the decision. They need to understand the probabilities of loss (i.e., risk) so they can balance those against the reward probabilities.

Risk tolerance

One of the questions that FAIR and other risk analysis methods will never answer is whether a given level of risk is acceptable. Risk analysis only identifies how much risk exists. We can draw lines on charts and establish various criteria that attempt to delineate between what's acceptable and unacceptable but, at the end of the day, acceptability is a very human and personal issue. A decision-maker always chooses between risk and reward, or between various risks. That's what risk decisions are – choices between the probability of loss (risk) and the probability of reward. Risk analysis provides only half of the information needed for the decision.

Another thing to keep in mind is that risk tolerance is unique to every individual. We each have different tolerances for loss, and our tolerance for loss varies from issue to issue. For example, I may have a very low tolerance for financial loss, but be entirely willing to take up skydiving. As a result, we shouldn't become too concerned when others have a very different perspective on what represents acceptable risk. Those differences are normal, natural, and unavoidable.

Summary

Everything I've covered so far highlights the fact that information risk is a complex subject, and that our profession has been challenged to deal with it effectively. At this point I'll add that FAIR is not a perfect solution; there are no perfect solutions. FAIR does, however, provide a rational, effective, and defensible solution to the challenges I've described.

Risk Landscape Components

Before we can begin to analyze any risk scenario, we have to understand the components that make up the landscape. The FAIR framework contains four primary components – threats, assets, the organization itself, and the external environment. Everything within a scenario falls into one of these categories, and each has attributes, or factors, that contribute positively or negatively to risk.

In this section, we'll spend most of our time covering the threat component because even simple FAIR analyses rely on the analyst to have a solid understanding of threat concepts.

Threats

As I mentioned in the Bald Tire section, threats are anything (e.g., object, substance, human, etc.) that are capable of acting against an asset in a manner that can result in harm. A tornado is a threat, as is a flood, as is a hacker. The key consideration is that threats apply the force (water, wind, exploit code, etc.) against an asset that can cause a loss event to occur.

As we progress through this document, we'll see that threat factors play a major role in our loss probabilities. The challenge is that we can't know who the next attacker will be any more than we can know whether the next toss of the coin will turn up heads. However, because we understand the fundamental characteristics of the coin and the toss, we can reasonably predict that out of the next 500 tosses, about (but probably not precisely) 50% will turn up heads. In the same way, we can define and characterize the threat landscape, and then establish reasoned probabilities regarding the frequency and nature of attacks.

Threat Agents

Individuals within a threat population

Practically anyone and anything can, under the right circumstances, be a threat agent – the well-intentioned, but inept, computer operator who trashes a daily batch job by typing the wrong command, the regulator performing an audit, or the squirrel that chews through a data cable.

Threat Communities

Subsets of the overall threat agent population that share key characteristics

The notion of threat communities is a powerful tool for understanding who and what we're up against as we try to manage risk. For example, consider the following threat community profile:

- ! Motive: ideology
- ! Primary intent: damage/destroy
- ! Sponsorship: unofficial
- ! Preferred general target characteristics: entities or people who clearly represent a conflicting ideology
- ! Preferred specific target characteristics: high profile, high visibility
- ! Preferred targets: human, infrastructure (buildings, communications, power, etc.)
- ! Capability: varies by attack vector (technological: moderate)
- ! Personal risk tolerance: high
- ! Concern for collateral damage: low

A threat agent having these characteristics might be said to fall into the Terrorist threat community.

The probability that your organization would be subject to an attack from the terrorist threat community would depend in large part on the characteristics of your organization relative to the motives, intents, and capabilities of the terrorists. Is your organization closely affiliated with ideology that conflicts with known, active terrorist groups? Does your organization represent a high profile, high impact target? Is your organization a soft target? How does your organization compare with other potential targets? If your organization were to come under attack, what components of your organization would be likely targets? For example, how likely is it that terrorists would target your information or systems?

The following threat communities are examples of the human malicious threat landscape many organizations face:

- ! Internal
 - Employees
 - Contractors (and vendors)
 - Partners
- ! External
 - Cyber-criminals (professional hackers)
 - Spies
 - Non-professional hackers

- Activists
- Nation-state intelligence services (e.g., counterparts to the CIA, etc.)
- Malware (virus/worm/etc.) authors

Note that you can subdivide the threat population further, or differently, as suits your needs. For example, in many risk analyses it makes perfect sense to subdivide employees into those who have elevated access privileges and greater technical expertise (e.g., system and network administrators), and those who don't have elevated privileges or high levels of expertise (e.g., the general employee populace). When you subdivide communities or identify new communities, it's important to be clear on what differentiates the new communities from existing ones.

It's also important to recognize that threat community membership isn't mutually exclusive. In other words, a threat agent can be a member of more than one threat community – e.g., a non-professional hacker might also be an employee or contractor. Similarly, the characteristics of individual threat agents may not always align perfectly with any single threat community. In other words, the characteristics of an individual threat agent may not align with every characteristic of the terrorist community. You might, for example, have a “terrorist” with a low tolerance for personal risk. Remember, the point isn't to develop a perfect characterization of the threat landscape, as that's not possible. The point is to develop a reasoned and more thorough understanding of the threat landscape. This allows us to better estimate probabilities and identify more effective risk management solutions.

Threat Characteristics

We can identify any number and variety of threat agent characteristics with which to profile threat communities. Under most circumstances there are relatively few truly significant characteristics. Including too many characteristics in our analysis makes the model much more difficult to use, with relatively little improvement in results. This is an example of where risk modeling typically will trade precision for increased practicality.

There are four primary components of our risk taxonomy that we want to identify threat agent characteristics for – those characteristics that affect:

- ! The frequency with which threat agents come into contact with our organizations or assets
- ! The probability that threat agents will act against our organizations or assets
- ! The probability of threat agent actions being successful in overcoming protective controls
- ! The probable nature (type and severity) of impact to our assets

It's important for us to understand the factors that drive these differentiating characteristics in order to effectively assess the probability of being subject to attack and, if subjected to attack, the likely nature, objective, and outcome of the attack. We'll examine these factors a bit more as we go along.

Assets

Within the information risk landscape, we can define *Asset* as any data, device, or other component of the environment that supports information-related activities, and which can be affected in a manner that results in loss. Assets have characteristics related to value, liability, and controls strength that represent risk factors.

In order for an asset to introduce any potential for loss, it has to have one or more characteristics that represent value or liability. For example, an organization's productivity has to depend on an asset before harm to that asset can result in productivity loss. Likewise, regardless of the sensitivity of an asset, an organization has to have a legal duty to protect the asset in order for the asset to represent a potential legal liability.

For this introduction to FAIR, we'll limit our asset value and liability considerations to:

- ! **Criticality** – that characteristic of an asset that has to do with the impact to an organization's productivity. For example, the impact a corrupted database would have on the organization's ability to generate revenue
- ! **Cost** – the costs associated with replacing an asset that has been stolen or destroyed. Examples include the cost of replacing a stolen laptop or rebuilding a bombed-out building
- ! **Sensitivity** – the impact resulting from confidential information being disclosed or improperly used

The Organization

Risk exists within the context of an organization or entity. In other words, harm to assets affects one or more of the organization's value propositions (more on this later). It is the organization that loses resources or the ability to operate. Characteristics of the organization also can serve to attract the attention of certain threat communities, which may increase the frequency of events.

The External Environment

The environment in which an organization operates plays a significant role in risk. Various external characteristics, such as the regulatory landscape, competition within the industry, etc., all help to drive the probability of loss.

We'll cover organizational and external environment factors in greater detail in the *Factoring* and *Measurement* sections.

Risk Factoring

In this section, we'll begin to define and develop a taxonomy for information risk by decomposing it into its fundamental components. The resulting structure provides a solid foundation for risk analysis and an effective basis for explaining analysis results.

Note that we won't go into great detail on any of the factors, nor will we discuss measuring the various factors. We'll cover measurement in a later section.

Decomposing Risk

Earlier in this document, I provided the following definition for risk:

The probable frequency and probable magnitude of future loss




Figure 1

With this as a starting point, the first two obvious components of risk are loss frequency and loss magnitude. In FAIR, these are referred to as *Loss Event Frequency* (LEF) and *Probable Loss Magnitude* (PLM), respectively.

We'll decompose the factors that drive loss event frequency first, and then examine the factors that drive loss magnitude.

Loss Event Frequency (LEF)

The probable frequency, within a given timeframe, that a threat agent will inflict harm upon an asset.

In order for a loss event to occur, a threat agent has to act against an asset, and that action has to result in loss. This leads us to our next two factors: *Threat Event Frequency* (TEF) and *Vulnerability*.




Figure 2

Note that time-framing is key to differentiating between possibility and probability because, given enough time, almost any event is possible. By time-framing our analysis, we're more or less forced to treat the issue as a probability.

Threat Event Frequency (TEF)

The probable frequency, within a given timeframe, that a threat agent will act against an asset.

You probably see the similarity between this definition and the definition for LEF. The only difference is that the definition for Threat Event Frequency doesn't include whether threat agent actions are successful. In other words, threat agents may act against assets, but be unsuccessful in affecting the asset. A common example would be the hacker who unsuccessfully attacks a web server. Such an attack would be considered a threat event, but not a loss event.

This definition also provides us with the two factors that drive threat event frequency; *Contact* and *Action*. Note that action is predicated upon contact. Figure 3 adds these two factors to our taxonomy.



Figure 3

Contact

The probable frequency, within a given timeframe, that a threat agent will come into contact with an asset.

Contact can be physical or “logical” (e.g., over the network). Regardless of contact mode, three types of contact can take place; random, regular, and intentional.

- ! **Random** – the threat agent “stumbles upon” the asset during the course of unfocused or undirected activity
- ! **Regular** – contact occurs because of the regular actions of the threat agent. For example, if the cleaning crew regularly comes by at 5:15, leaving cash on top of the desk during that timeframe sets the stage for contact
- ! **Intentional** – the threat agent seeks out specific targets

Each of these types of contact is driven by various factors. Because this is only an introduction, we won’t get into the details at this time. A useful analogy, however, is to consider a container of fluid containing two types of suspended particles – threat particles and asset particles. The probability of contact between members of these two sets of particles is driven by various factors, including:

- ! Size (surface area) of the particles
- ! The number of particles
- ! Volume of the container
- ! How active the particles are
- ! Viscosity of the fluid
- ! Whether particles are attracted to one another in some fashion
- ! Etc...

Action

The probability that a threat agent will act against an asset once contact occurs.

Once contact occurs between a threat agent and an asset, action against the asset may or may not take place. For some threat agent types, action always takes place. For example, if a tornado comes into contact with a house, action is a foregone conclusion. Action is only in question when we're talking about "thinking" threat agents such as humans and other animals, and artificially intelligent threat agents like malicious programs (which are extensions of their human creators).

The probability that an intentional malicious act will take place is driven by three primary factors:

- ! **Asset value** – from the threat agent's perspective
- ! **Level of effort** – the threat agent's expectation of how much effort it will take to compromise the asset
- ! **Risk** – the probability of negative consequences to the threat agent – i.e., the probability of getting caught and suffering unacceptable consequences.

Vulnerability

Having covered the high-level factors that drive whether threat events take place, we now turn our attention to the factors that drive whether the asset is able to resist threat agent actions. Vulnerability is defined as:

The probability that an asset will be unable to resist the actions of a threat agent.

As you'll recall from the Introduction, vulnerability exists when there's a difference between the force being applied by the threat agent, and an object's ability to resist that force. This simple analysis provides us with the two primary factors that drive vulnerability; *Threat Capability* and *Control Strength*. Figure 4, below, adds these factors to our taxonomy.

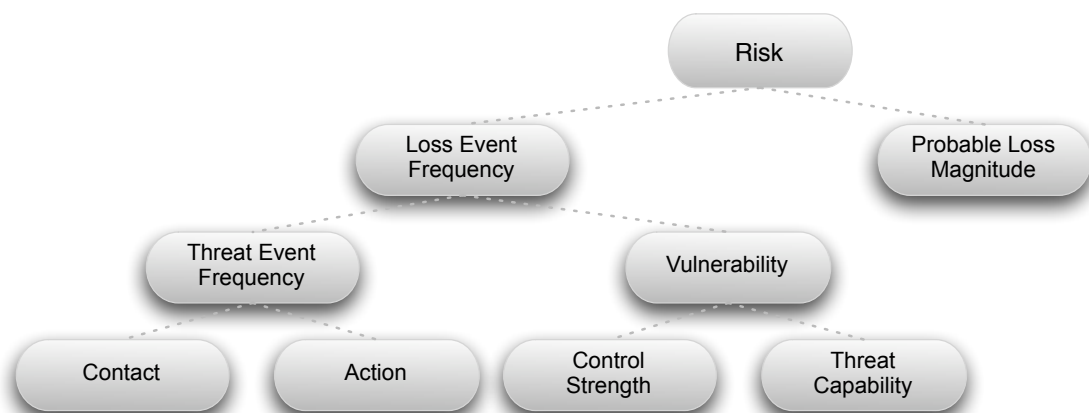


Figure 4

Vulnerability is always relative to the type of force involved. In other words, the tensile strength of a rope is pertinent only if the threat agent force is a weight applied along the length of the rope. Tensile strength doesn't generally apply to a scenario where the threat agent is fire, chemical erosion, etc. Likewise, an antivirus product doesn't provide much in the way of protection from the internal employee seeking to perpetrate fraud. The key, then, is to evaluate vulnerability in the context of specific threat types and control types.

One final point regarding vulnerability – there's no such thing as being more than 100% vulnerable to any specific threat agent/attack vector combination. Vulnerability can exist such that harm can occur from more than one threat agent through more than one attack vector, but each of those represents a different potential threat event. For example, if I'm walking down the street at night in a particularly dangerous part of town, I'm vulnerable to multiple potential threat events, for example – being run over by a car, being mugged, or being the victim of a drive-by shooting. My vulnerability to any one of these events cannot exceed 100%, yet my aggregate risk is obviously greater as a result of the multiple threat scenarios.

Threat Capability

The probable level of force that a threat agent is capable of applying against an asset.

Not all threat agents are created equal. In fact, threat agents within a single threat community are not all going to have the same capabilities. What this should tell us is that the probability of the most capable threat agent acting against an asset is something less than 100%. In fact, depending upon the threat community under analysis and other conditions within the scenario, the probability of encountering a highly capable threat agent may be remote.

As information security professionals we often struggle with the notion of considering threat agent capability as a probability. We tend, instead, to gravitate toward focusing on the worst case. But if we look closely at the issue, it's clear that focusing solely on worst case is to think in terms of possibility rather than probability.

Another important consideration is that some threat agents may be very proficient in applying one type of force, and incompetent at others. For example, a network engineer is likely to be proficient at applying technological forms of attack, but may be relatively incapable of executing complex accounting fraud.

Control Strength

The strength of a control as compared to a baseline measure of force.

Clear as mud? Let's look at an example...

A rope's tensile strength rating provides an indication of how much force it's capable of resisting. The baseline measure (CS) for this rating is pounds per square inch (PSI), which is determined by the rope's design and construction. This CS rating doesn't change when the rope is put to use. Regardless of whether you have a 10-pound weight on the end of that 500-PSI rope, or a 2000-pound weight, the CS doesn't change.

Unfortunately, the information risk realm doesn't have a baseline scale for force that's as well defined as PSI. A deeper discussion of how we can deal with this will come later. For now, consider password strength as a simple example. We can estimate that a password eight characters long, comprised of a mixture of upper and lower case letters, numbers, and special characters will resist the cracking attempts of some percentage of the general threat agent population. The password control strength (CS) can be represented as this percentage. (Recall that CS is relative to a particular type of force – in this case cracking.) Vulnerability is determined by comparing CS against the capability of the specific threat community under analysis. For example, password CS may be estimated at 80%, yet the threat community within a scenario might be estimated to have better than average capabilities – let's say in the 90% range. The difference represents vulnerability.

Clearly, reality is more complex than this, and a full discussion of FAIR will address this greater complexity. For now, we'll keep things relatively simple.

Probable Loss Magnitude (PLM)

The previous section introduced the factors that drive the probability of loss events occurring. This section describes the other half of the risk equation – the factors that drive loss magnitude when events occur.

Unfortunately, loss is one of the toughest nuts to crack in analyzing risk. Various approaches have been tried, with varying degrees of success, but none have gained widespread use or acceptance. As a result, we often exclude loss considerations altogether, we only cite the worst-case possibilities, or we try to be precise in our calculations. Excluding loss from an analysis means that we aren't analyzing risk (by definition, risk always has a loss component). Citing worst-case possibilities alone removes the probabilistic element from our analysis (by definition, risk is a probability issue). Trying to be precise is generally a waste of time because of the inherent complexity within loss, and because decision-makers generally only need a ballpark idea of the loss probabilities. Their experience with other forms of risk (investment, market, etc.) has taught them that actual losses can't be predicted with any precision.

There are a number of reasons why it's difficult to evaluate loss probability, for example:

- ! It's very difficult to put a precise value on assets at risk
- ! Assets generally have more than one value or liability characteristic
- ! Loss can take many forms
- ! A single event can result in more than one form of loss
- ! Complex systemic relationships exist between the different forms of loss
- ! Many factors determine loss magnitude

Making matters even more difficult in the information risk environment is the fact that we have very little good data regarding loss magnitude. Many organizations don't perform loss analysis when events occur, and those that do track loss often limit their analyses to the 'easy stuff' (e.g., person-hours, equipment replacement, etc.). Furthermore, without a standard taxonomy it's very difficult to normalize the data across organizations.

Before we go any farther, my experience has been that loss from information security incidents generally has a distribution that looks something like the following:




Figure 5

In other words, there are far more events that result in loss at the low end of the magnitude spectrum than there are at the high end of the spectrum. For example, individual virus incidents, unauthorized use of systems to serve up MP3 files, even password cracking and web site defacement, rarely result in significant loss. The question we have to ask ourselves is, “Why?” What factors are responsible for this? Clearly some of these events have significant potential for harm, but if we compared the actual loss from two similar events – one in which minimal loss occurred, and another where substantial loss occurred – what factors determined the difference? In order for us to make reasoned estimates of loss, we have to understand these factors.

Forms of Loss

An asset’s loss potential stems from the value it represents and/or the liability it introduces to an organization. For example, customer information provides value through its role in generating revenue for a commercial organization. That same information also can introduce liability to the organization if a legal duty exists to protect it, or if customers have an expectation that the information about them will be appropriately protected.

Six forms of loss are defined within FAIR – productivity, response, replacement, fines/judgments (F/J), competitive advantage (CA), and reputation.

- ! **Productivity** – the reduction in an organization’s ability to generate its primary value proposition (e.g., income, goods, services, etc.)
- ! **Response** – expenses associated with managing a loss event (e.g., internal or external person-hours, logistical expenses, etc.)
- ! **Replacement** – the intrinsic value of an asset. Typically represented as the capital expense associated with replacing lost or damaged assets (e.g., rebuilding a facility, purchasing a replacement laptop, etc.)
- ! **Fines and judgments (F/J)** – legal or regulatory actions levied against an organization. Note that this includes bail for any organization members who are arrested.
- ! **Competitive advantage (CA)** – losses associated with diminished competitive advantage. Within this framework, CA loss is specifically associated with assets that provide competitive differentiation between the

organization and its competition. Within the commercial world, examples would include trade secrets, merger and acquisition plans, etc. Outside of the commercial world, examples would include military secrets, secret alliances, etc.

! **Reputation** – losses associated with an external perception that an organization's leadership is incompetent, criminal, or unethical

Keep in mind that loss is always evaluated from a single perspective – typically that of the organization under analysis. For example, although customers might be harmed if their personal information is stolen, our risk analysis would evaluate the losses experienced by the organization rather than the losses experienced by the customers.

Loss Factors

All loss factors fall within one of the following four categories – asset, threat, organization, and external. For reasons that will become clear as we go along, asset and threat loss factors are referred to as primary loss factors, while organizational and external loss factors are referred to as secondary loss factors.

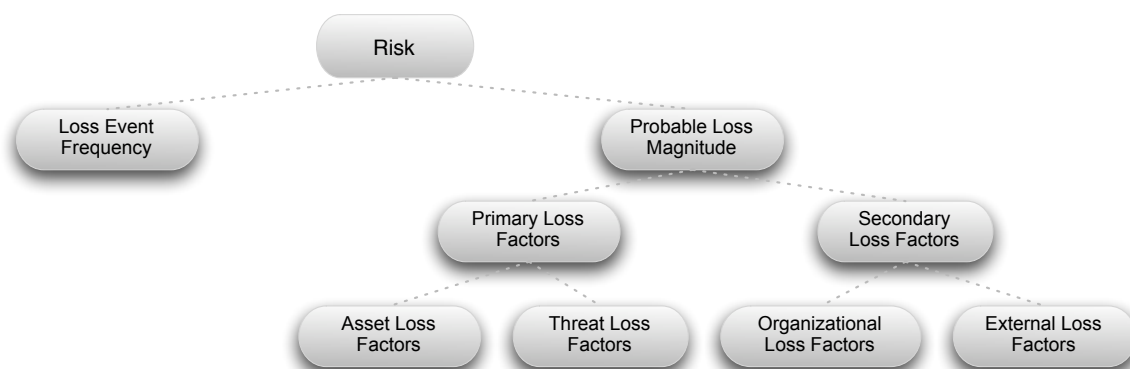


Figure 6

In order for us to make reasoned judgments about the form and magnitude of loss within any given scenario, we have to evaluate the factors within all four of these categories. Within this introduction, we'll limit our discussion to some of the most common and most important loss factors.

Primary Loss Factors

Asset loss factors

There are two asset loss factors that we are concerned with – value/liability, and volume.

As I alluded to above, and as we'll see when we cover measurement, the *value/liability* characteristics of an asset play a key role in both the nature and magnitude of loss. We can further define value/liability as:

- ! **Criticality** – characteristics of an asset that have to do with the impact to an organization's productivity. For example, the impact a corrupted database would have on the organization's ability to generate revenue
- ! **Cost** – refers to the intrinsic value of the asset – i.e., the cost associated with replacing it if it's been made unavailable (e.g., stolen, destroyed, etc.). Examples include the cost of replacing a stolen laptop or rebuilding a bombed-out building
- ! **Sensitivity** – the harm that can occur from unintended disclosure. Sensitivity is further broken down into four sub-categories:
 - **Embarrassment/reputation** – the information provides evidence of incompetent, criminal, or unethical management. Note that this refers to reputation damage resulting from the nature of the information itself, as opposed to reputation damage that may result when a loss event takes place.
 - **Competitive advantage** – the information provides competitive advantage (e.g., key strategies, trade secrets, etc.). Of the sensitivity categories, this is the only one where the sensitivity represents value. In all other cases, sensitivity represents liability.
 - **Legal/regulatory** – the organization is bound by law to protect the information
 - **General** – sensitive information that doesn't fall into any of the above categories, but would result in some form of loss if disclosed

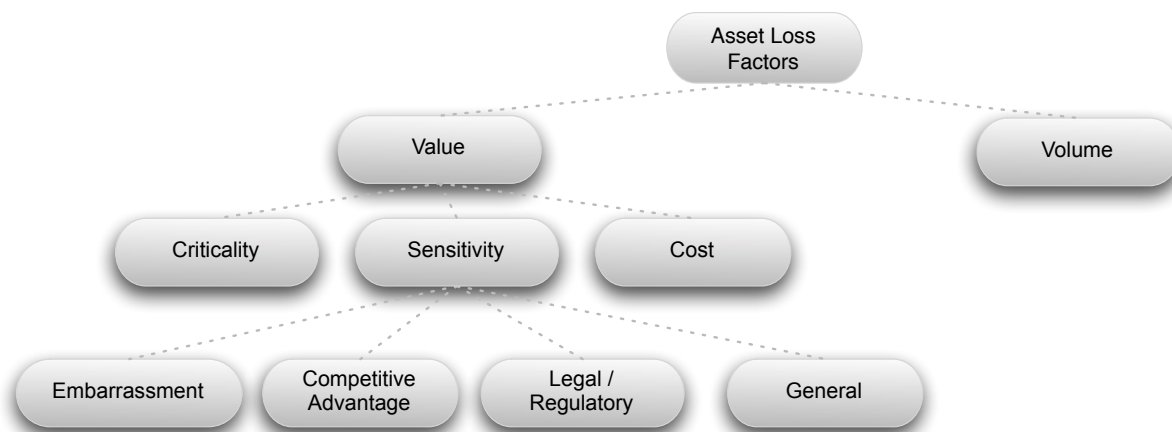


Figure 7

Asset volume simply recognizes that more assets at risk equal greater loss magnitude if an event occurs – e.g., two children on a rope swing versus one child, or one sensitive customer record versus a thousand.

Threat loss factors

Within this document, we'll limit our threat considerations to three loss factors – action, competence, and whether the threat agent is internal or external to the organization.

Threat agents can take one or more of the following actions against an asset:

- ! **Access** – simple unauthorized access
- ! **Misuse** – unauthorized use of assets (e.g., identity theft, setting up a porn distribution service on a compromised server, etc.)
- ! **Disclose** – the threat agent illicitly discloses sensitive information
- ! **Modify** – unauthorized changes to an asset
- ! **Deny access** – includes destruction, theft of a non-data asset, etc.

It's important to recognize that each of these actions affects different assets differently, which drives the degree and nature of loss. For example, the potential for productivity loss resulting from a destroyed or stolen asset depends upon how critical that asset is to the organization's productivity. If a critical asset is simply illicitly accessed, there is no direct productivity loss. Similarly, the destruction of a highly sensitive asset that doesn't play a critical role in productivity won't directly result in a significant productivity loss. Yet that same asset, if disclosed, can result in significant loss of competitive advantage or reputation, and generate legal costs. The point is that it's the combination of the asset and type of action against the asset that determines the fundamental nature and degree of loss.

Which action(s) a threat agent takes will be driven primarily by that agent's motive (e.g., financial gain, revenge, recreation, etc.) and the nature of the asset. For example, a threat agent bent on financial gain is less likely to destroy a critical server than they are to steal an easily pawned asset like a laptop. As we'll see when we cover measurement, it's critical to have a clear definition of your threat community in order to effectively evaluate loss magnitude.

Threat Competence is similar to the Threat Capability factor that contributes to vulnerability. The difference is subtle, but important. Threat Competence has to do with the amount of damage a threat agent is able to inflict, while Threat Capability has to do with the threat agent's ability to put itself in a position to inflict harm. An example may help to clarify this point. A terrorist threat agent has capabilities they would employ in an attempt to access nuclear secrets. These capabilities play a role in the likelihood that they'll be successful in gaining access. Their ability to inflict harm once they've acquired the secrets (e.g., build a bomb) is, however, dependent upon a different set of competencies. In FAIR, the characteristics that enable the terrorist to compromise defenses and be in a position to acquire the secrets are called Threat Capabilities. The characteristics that enable them to inflict harm (e.g., create a bomb) are referred to as Threat Competence. We won't dwell on Threat Competence in this document. Nonetheless, it's useful to recognize that this factor exists in order to have a more complete understanding of risk.

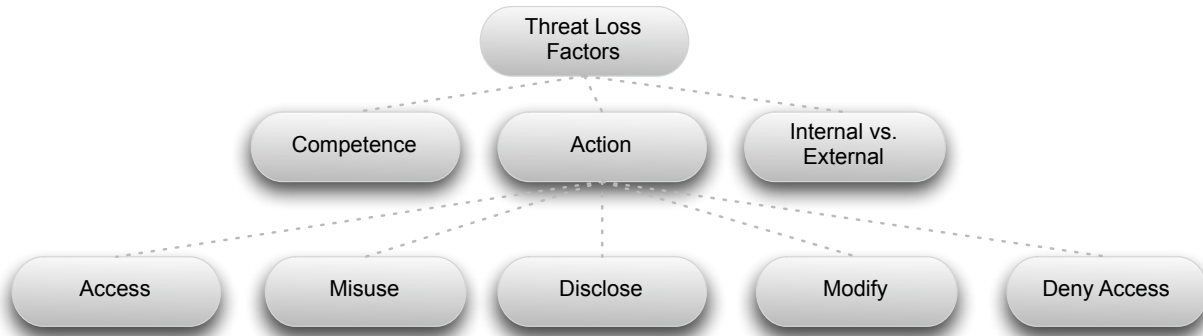


Figure 8

The consideration of whether a threat agent is external or internal to the organization can play a pivotal role in how much loss occurs. Specifically, loss events generated by malicious internal threat agents (including employees, contractors, etc.) typically have not resulted in significant regulatory or reputational losses because it's recognized that trusted insiders are exceedingly difficult to protect against. We'll see in the next section that other factors contribute to how much leeway an organization is granted from insider events.

Secondary Loss Factors

Secondary loss factors are those organizational and external characteristics of the environment that influence the nature and degree of loss.

Organizational Loss Factors

There are many organizational loss factors. Within this document, we'll limit our discussion to four – timing, due diligence, response, and detection.

The *timing* of an event can have a tremendous impact on loss. For example, an event occurring in the midst of a big advertising campaign may create significantly greater loss than a similar event at some other time of year.

Due *diligence* can play a significant role in the degree of liability an organization faces from an event. If reasonable preventative measures were not in place (given the threat environment and value of the asset), then legal and reputational damage can be far more severe. The challenge is that 'reasonable preventative measures' are not universally defined or agreed upon. Often, 'industry standards' or theoretical 'best practices' are looked to as guidelines for due diligence. Unfortunately, these guidelines typically don't consider the threat environment or loss magnitude. Consequently, industry standards and best practices may be insufficient (i.e., not truly representative of due diligence) or overly conservative (i.e., prohibitively expensive given the real risk).

How effectively an organization *responds* to an event can spell the difference between an event nobody remembers a year later, and one that stands out as an example (good or bad) in the annals of history. There are three components to response:

- ! **Containment** – has to do with an organization’s ability to limit the breadth and depth of an event – for example, cordoning-off the network to contain the spread of a worm
- ! **Remediation** – has to do with an organization’s ability to remove the threat agent – e.g., eradicating the worm
- ! **Recovery** – refers to the ability to bring things back to normal

All three of these response components must exist, and the degree to which any of them is deficient can have a significant impact on loss magnitude.

We tend to think of response capabilities solely within the context of criticality – i.e., the ability to return productivity to normal. It’s critical to recognize, however, that response capabilities also can significantly affect losses resulting from sensitive information disclosure. For example, an organization that experiences a publicly disclosed breach of confidential customer information generally can significantly reduce its losses by being forthright in its admissions, and by fully compensating harmed parties. Conversely, an organization that denies and deflects responsibility is much more likely to become a pariah and a media whipping post.

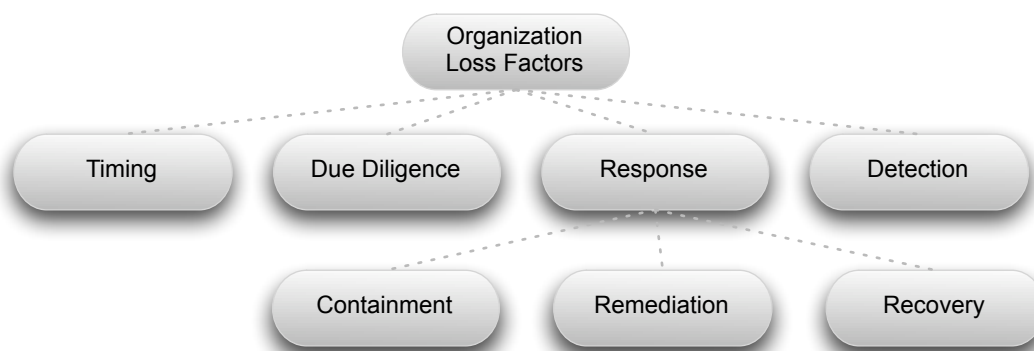


Figure 9

You can’t respond to something you haven’t detected – i.e., response is predicated on detection. In training sessions, the question often comes up, “What about those events we may not know about – the corporate spies, etc.?” Clearly, incidents take place that don’t show up on the radar. However, it’s also reasonable to believe that such events – if they result in material loss – will almost always be detected eventually. For example, the damage from sensitive competitive advantage information that makes its way to a competitor will materialize and almost certainly be recognized. Was the detection timely? Perhaps not. However, once detected, the organization still may have an opportunity to respond and reduce its losses. For example, legal action against a competitor who stole proprietary information might be appropriate. The point is that material loss is almost certain to be detected, and with detection comes an opportunity to respond and manage loss magnitude.

External Loss Factors

External loss factors generally fall into one of the following four categories – the legal and regulatory landscape, the competitive landscape, the media, and external stakeholders (e.g., customers, partners, stockholders, etc.). A couple of important things to recognize about external loss factors include:

- ! These four categories represent entities that can inflict a secondary form of harm upon the organization as a consequence of an event. In other words, events will often result in direct forms of loss (e.g., productivity, response, replacement) due to the criticality and inherent value characteristics of assets. Secondary losses may also occur based upon the external reaction to a loss event (e.g., sensitive information disclosure, etc.).
- ! All of the factors within these external categories can be described as “reactive to an event.” In other words, in order for an external factor to affect loss magnitude, the event has to be detected by an external entity. For example, if an employee executes identity theft by misusing their legitimate access to customer information, the customer(s), regulators, and lawyers can’t inflict harm upon the organization unless the identity theft is tied back to the organization. Likewise, if a productivity outage occurs but isn’t detected by customers, partners, etc., then the organization will not be subject to a negative response on the part of those stakeholders.

This last point leads us to our first external loss factor – *detection*. Based upon the premise above, we can think of detection as a binary factor that all other external factors are predicated on. External detection of an event can happen as a consequence of the severity of the event, through intentional actions by the threat agent, through unauthorized disclosure by someone on the inside who’s familiar with the event, intentional disclosure by the organization (either out of sense of duty, or because it’s required by law), or by accident.

The regulatory and legal landscape is primarily made up of three parts – regulations (local, state, federal and international), contract law, and case law. I’m not a lawyer, and am not qualified to discuss these issues at length. Furthermore, this component of the external landscape is evolving rapidly, and so any conclusions I might draw here could be outdated by the time you read this. For now, suffice it to say that fines and sanctions can be significant for organizations within regulated industries. In theory, however, fines and judgments are driven in part by how much harm actually occurs from an event and the level of due diligence exercised to prevent it from occurring in the first place. In other words, if an event occurs that represents a regulatory or legal breach, fines and judgments should reflect how much harm actually occurs to the affected stakeholders as well as how proactive the organization was in preventing the loss.

Losses associated with the competitive landscape typically have to do with the competition’s ability to take advantage of the situation. For example, if an organization experiences an event that causes its stakeholders to consider taking their business elsewhere, a competitor’s ability to leverage that weakness will affect how much loss occurs.

Media reaction can have a powerful affect on how stakeholders, lawyers, and even regulators and competitors view the event. If the media chooses to vilify the organization, and keep it on the headlines for an extended period, the result can be devastating. Conversely, if the media paints the organization as a well-intentioned victim who exercised due diligence but still suffered the event at the hands of a criminal, then legal and reputation damage can be minimized. This is why organizations must have effective crisis communication processes in place.

External stakeholders generally inflict harm by taking their business elsewhere – i.e., supporting a rival. This happens when they:

- ! Have been harmed directly as a result of an incident. The organization's response to the event is crucial in mitigating this exposure
- ! Perceive that their interests are better served elsewhere (the organization's value proposition is diminished). Here again, an organization generally has some opportunity to mitigate this exposure through prompt, effective action
- ! View the organization (or, more accurately, its leadership) as incompetent, untrustworthy, and/or criminal. This can be a much tougher exposure to mitigate




Figure 10

As you can see, the complex nature of loss doesn't lend itself to simple analysis. Rest assured, we won't measure every discreet factor within a scenario. Instead, we'll perform our evaluation at a higher level of abstraction, but with the benefit of knowing what the contributing factors are that exist within the deeper layers of the model.

Controls

Because our profession has tended to focus on controls, much of the information in this section should be familiar to most of you. That said, there's more complexity to controls than meets the eye, and a separate whitepaper is warranted in order to cover the topic thoroughly. Nonetheless, we'll cover some of the higher-level concepts in order to flesh-out the framework. Specifically, this section will introduce:

- ! Three control dimensions:

- Form
- Purpose
- Category

- ! The control lifecycle

Control dimensions

All controls can be characterized through three dimensions. These characterizations can help us understand where a control fits within the risk framework, enables us to better assess control capabilities, and helps us eliminate gaps in our organization's risk management program.

Forms

Controls take one of three forms – policy, process, or technology. An important thing to keep in mind is that relatively few controls stand alone. For example, a technology control, such as a firewall, will have policies associated with how it

should be configured and used, and processes that determine how it's managed. Likewise, the firewall can be considered an instantiation of a “least privilege” policy control. Bottom line – controls interact through a complex set of dependencies. By understanding these systemic interdependencies, we can identify levers that enable us to have greater effect with less effort.

Purpose

Control purpose refers to whether the controls are primarily preventive, detective, or responsive in nature. The qualifier “primarily” is important to keep in mind, as some controls may serve more than one purpose – killing more than one bird with one stone. Antivirus software is a terrific example, as it serves all three purposes to some degree. This multi-purpose characteristic is important when we evaluate potential control solutions, as it's usually to our advantage when it's available.

Categories

Unlike form and purpose, the *control category* dimension may not be familiar to you. This dimension provides an explicit taxonomy for controls, which helps to ensure that gaps don't exist in your controls environment.

There are three primary control categories:

- ! Loss event controls
- ! Threat event controls
- ! Vulnerability controls

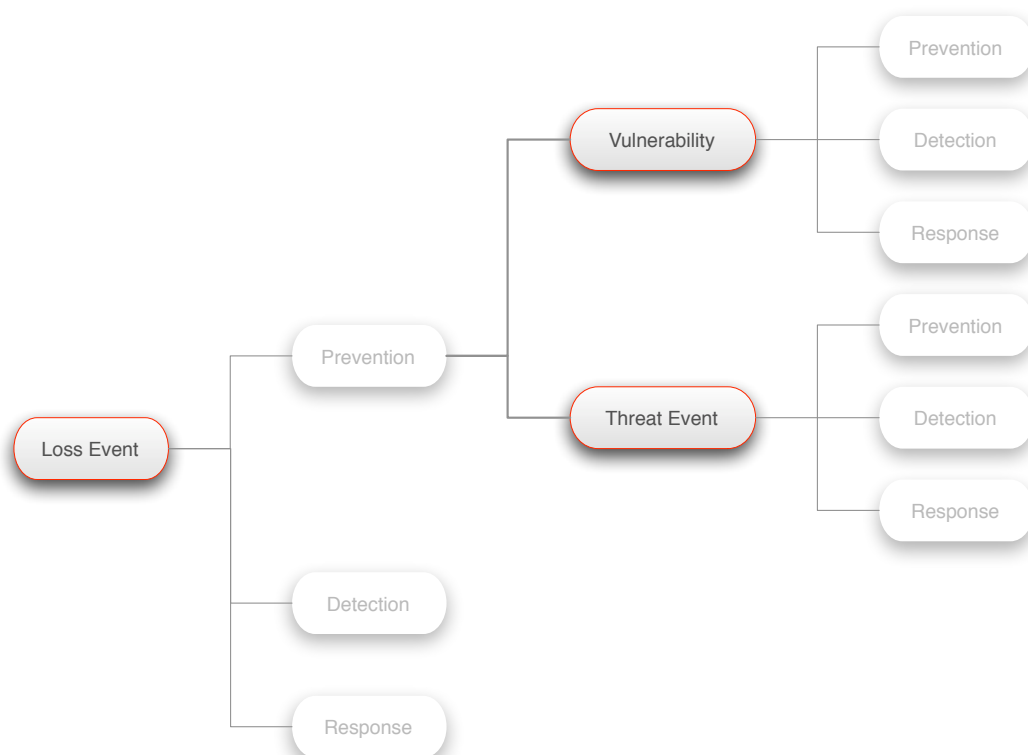


Figure 11

Each of the preventative, detective, and response controls we're familiar with, fall into one of these categories. Examples include:

- ! **Vulnerability/Prevention controls** consist of measures that make compromise more difficult. The most common example would be authentication.
- ! **Vulnerability/Detection controls** identify existing vulnerabilities. Attack and penetration exercises and system scans are examples.
- ! **Loss Event/Response controls** help to reduce the magnitude of loss when an event occurs. Disaster recovery processes and fail-over systems are examples of this type of control.
- ! **Threat Event/Prevention controls** primarily are responsible for limiting threat agent contact or action. A common example would be a multilayer Internet DMZ architecture.

As I mentioned earlier, some controls play a role in more than one control category. In this introduction to FAIR, and in the risk analysis scenario presented later on, most of our attention will be on threat event/prevention controls.

Control Lifecycle

All controls have a four-stage lifecycle – design, implementation, use/maintenance, and disposal. We won't delve into these stages at this time, but recognize that vulnerability can and often does fluctuate or increase over time due to either of two conditions:

- ! Changes that occur to the control itself, and/or
- ! Changes within the threat landscape

For example, during maintenance, a systems administrator can intentionally or inadvertently reduce a server's control strength by changing a system parameter. In this case, a change has occurred to the control itself. Likewise, vulnerability can be increased if a new capability is introduced into the threat population that potentially would allow an attacker to bypass or defeat the control (e.g., a new password cracking tool). In this case, the threat landscape has changed.

Because of this variability, it's critical that we have the ability to detect and respond in a timely manner to either type of change in order to maintain the desired level of vulnerability. The latency between when a change occurs, and when we detect and mitigate the change, often accounts for a significant portion of our overall vulnerability.

Measuring Risk

This section will begin to answer the reasonable question, “How can we credibly measure the various risk factors?” Before we answer this question, we need to cover some fundamental measurement concepts.

Measurement Theory

Subjective vs. objective measures

A commonly expressed concern has been – “*Without data we’re relying upon subjectivity!*” This is true, at least to some degree. However, subjectivity is also a part of every other analysis method known to man – even those that benefit from reams of data. Who do you suppose establishes the criteria for data collection, collects and filters the data, processes the data, and draws conclusions from data? Eternally flawed, highly subjective, human beings do. My actuarial colleagues pointed this out to me when I first asked them to critique FAIR. The simple fact is that anything done by humans inherently includes some degree of subjectivity. Therefore, it isn’t a matter of whether subjectivity is involved, because that’s a foregone conclusion. It’s a matter of how much. Our goal, then, is to bring as much objectivity as possible into the process. The FAIR framework goes a long way toward driving more objectivity into an analysis, but there’s no possible way to eliminate subjectivity completely.

There always will be those who reject any quantitative analysis that’s not founded upon empirical data. So be it. That’s a religious war that’s raged for decades in various scientific, social, and philosophical domains, and it’s an argument that isn’t going to be won or lost on the basis of any claim I present. Again, the simple fact is, we don’t have much good data today and yet we do have to treat information risk as a probability if we expect to manage it effectively. As a result, it’s

hard to argue against any logical framework that reasonably addresses our problem and pushes the “subjectivity—>objectivity meter” toward the right, especially given our existing alternatives.

There's no question, however, that good data can be very helpful in drawing and supporting conclusions about what might happen in the future. One of the benefits of FAIR is that the taxonomy establishes a basis for data collection criteria, provides a context for categorizing the data, as well as a framework for analyzing the data. If anything, this should help our profession develop the data necessary to further support our analyses.

Quantitative vs. qualitative measures

Discussions about risk analysis often come to the issue of which is better, qualitative or quantitative analysis. Decision-makers generally prefer quantitative measures if the method is credible. However, in the absence of taxonomy, reasonable measurement scales, and a computational engine that effectively reflects the relationships between risk factors, a quantitative method is not likely to be credible. The decision to use quantitative or qualitative measures should be driven by two things:

- ! The needs of the decision-makers who will use the results of the analysis
- ! The credibility of the available methods

FAIR can be used to perform qualitative or quantitative analysis. The only difference is in the measurement scales used. In other words, if you measure the risk factors using a qualitative scale such as High, Medium, or Low, then your outcome will be qualitative. Note, however, that if your qualitative labels actually refer to quantitative ranges (e.g., “‘High’ is greater than ten”) then you’re using a quantitative scale – albeit one with limited precision. This is neither good, nor bad. It’s just something to be aware of.

For those who are reluctant to perform quantitative analysis without empirical data, I strongly recommend using FAIR qualitatively so that the taxonomy and modeling components of the framework can still be leveraged.

What about precision?

Every measurement is, to some degree, an estimate. Measurement precision will always be limited by the characteristics of the measuring tool or reference, and by the capabilities (and intent) of the measurer.

Many of us in the information security field come from technology and engineering backgrounds, and if you ask a technically inclined person to measure something, he or she is often going to want to give you an answer to within several decimal points of precision. Unfortunately, risk doesn’t work that way. It’s unrealistic to believe that we can credibly establish precise estimations of probability in something as complex as information risk. Although precision would be nice, it isn’t necessary (nor should it be expected) in risk analysis. Remember – risk analysis is not the same thing as foretelling the future. It’s about providing rational estimates of loss probabilities. The better we analyze and understand our subject, the more likely it is that our estimates are going to be reasonable and useful. There will always be some degree of uncertainty and imprecision.

Measurement error & sensitivity

One of the key criteria for any good measurement method is consistency. In other words, if I have several analysts evaluate the same scenario, I should get reasonably consistent results. Accomplishing this requires well-defined measurement scales and well-trained analysts. Even so, some degree of inconsistency is natural given that experience and judgment will always vary, and no two humans will evaluate a complex subject in exactly the same manner. This is as true for FAIR as it is for any other complex measurement activity. Therefore, the pertinent question isn't whether error or inconsistency exists, but whether the degree of error or inconsistency is acceptable. Because risk analysis is not a precise endeavor, qualified analysts – those who have been effectively trained in the methodology – should have no difficulty using FAIR to produce acceptable results.

Measuring Risk Factors

In this section, we'll review some of the challenges we face in trying to measure information risk factors. I'll also introduce the simplified measurement approach that's used throughout the remainder of this document. Please keep in mind that a strong understanding of risk fundamentals and factors is far more important than the specific method used to measure risk.

Measurement challenges

Measurement scales already exist for some risk factors – dollars, for example, or the frequency of events within a given timeframe. For other factors – such as a human's ability to compromise a computer – no established scales exist. Neither, therefore, do we have a readily obvious way of measuring an ability to resist that type of force. This difficulty doesn't alter the fact that we need to be able to estimate these factors as credibly as possible.

The following paragraphs describe a simplified approach to measuring FAIR factors. Appendix A provides a form that can be used for similar analyses.

Measuring threat event frequency (TEF)

Recall that TEF represents the number of times, within a given timeframe, that we expect a particular threat community to act against an asset. Keep in mind that this is different from simply coming into contact with the asset. For example, systems and network engineers/administrators may regularly come into contact with sensitive information during the normal course of their jobs. In fact, the very nature of their job requires that they continually come into contact with critical systems. Fortunately, significant unauthorized or malicious acts are relatively rare. It's these acts that we want to estimate.

The following scale provides a simple means of estimating the probable TEF within a year. You'll notice that it's not precise, and that it uses qualitative labels for quantitative ranges (arguably making the scale quantitative). There's nothing sacred or mystical about how this scale was developed, and you're free to create your own scale. That said, our profession would benefit significantly by adopting a standard measurement scale. If you choose to create your own set

of scales, change them only if you absolutely have to. Any changes will create inconsistency between analyses performed using one scale versus another, and inconsistency is something we need to avoid.

Rating	Description
Very High (VH)	> 100 times per year
High (H)	Between 10 and 100 times per year
Moderate (M)	Between 1 and 10 times per year
Low (L)	Between .1 and 1 times per year
Very Low (VL)	< .1 times per year (less than once every ten years)

Measuring threat capability (Tcap)

For some threat types, scales of force already exist – for example; weight, heat, pressure, etc. For other threat types, such as a human's ability to compromise a computer, no established scales exist. Furthermore, because capability is derived from a combination of imprecise factors (e.g., skill and resources), it isn't possible to define an absolute measure.

In the absence of an absolute measurement scale, the simplified FAIR approach uses ratios. The notion is that the capability of any threat population can be described as a distribution. Some portion of the threat population will be more capable than the rest, and some portion will be less capable. Think 'bell curve'. Is this approach going to precisely match the capabilities of any given threat population? No. Nonetheless, it's far more accurate than assuming that all threat agents within a population have equal capabilities.

The following scale provides a simple means of estimating the Tcap of a threat community. Those of you with a statistical background may recognize that the H and VH levels are approximately plus one and two standard deviations (respectively) from the mean, and the L and VL levels are minus one and two standard deviations. Here again, we're using qualitative labels to represent quantitative ranges.

Rating	Description
Very High (VH)	Top 2% when compared against the overall threat population
High (H)	Top 16% when compared against the overall threat population
Moderate (M)	Average skill and resources (between bottom 16% and top 16%)
Low (L)	Bottom 16% when compared against the overall threat population
Very Low (VL)	Bottom 2% when compared against the overall threat population

Measuring control strength (CS)

The strength of any preventative control has to be measured against a baseline level of force. Unfortunately, just as with Tcap, no well-established scale exists. The good news is that we can leverage the Tcap scale as a baseline reference.

To use this scale, we simply estimate the strength of the control as measured against a threat community of average capability. For example, if we were estimating the CS of a four character, alphabetic only, password (against a cracking attack), we probably would rate it as 'Low'. This estimation would be based upon the fact that simple cracking tools are readily available, and that only relatively inept attackers would either have no access or wouldn't be capable of using them.

Rating	Description
Very High (VH)	Protects against all but the top 2% of an avg. threat population
High (H)	Protects against all but the top 16% of an avg. threat population
Moderate (M)	Protects against the average threat agent
Low (L)	Only protects against bottom 16% of an avg. threat population
Very Low (VL)	Only protects against bottom 2% of an avg. threat population

Deriving Vulnerability

Determining vulnerability is simple once you've established your Tcap and CS. Recall from the Factoring section that vulnerability is the difference between the force that's likely to be applied, and the asset's ability to resist that force. Using the matrix below, simply find the Tcap along the left side of the matrix, and the CS along the bottom. Where they intersect determines vulnerability.

		Vulnerability				
Tcap	VH	VH	VH	VH	H	M
	H	VH	VH	H	M	L
	M	VH	H	M	L	VL
	L	H	M	L	VL	VL
	VL	M	L	VL	VL	VL
		VL	L	M	H	VH
		Control Strength				

Deriving Loss Event Frequency (LEF)

Similar to vulnerability, LEF is a derived value. LEF is determined by intersecting TEF and Vulnerability within the matrix below. Recall that vulnerability is always a percentage and that it's not possible to be more than 100% vulnerable. As a result, LEF will never exceed TEF.

		Loss Event Frequency				
TEF	VH	M	H	VH	VH	VH
	H	L	M	H	H	H
	M	VL	L	M	M	M
	L	VL	VL	L	L	L
	VL	VL	VL	VL	VL	VL
		VL	L	M	H	VH
		Vulnerability				

Measuring probable loss magnitude (PLM)

It's generally a waste of time to establish precise loss estimates. Most scenarios are far too complex to allow for precise estimates, and decision-makers typically don't need or expect precision. Consequently, all we need to provide are reasonable ballpark estimates. Within this document, we'll use the scale below to make our estimates.

Magnitude	Range Low End	Range High End
Severe (SV)	\$10,000,000	--
High (H)	\$1,000,000	\$9,999,999
Significant (Sg)	\$100,000	\$999,999
Moderate (M)	\$10,000	\$99,999
Low (L)	\$1,000	\$9,999
Very Low (VL)	\$0	\$999

Note that this scale is ‘tuned’ for a specific organization’s capacity for loss. Organizations of different sizes or with fundamentally different (e.g., non-financial) value/loss propositions will want to create their own scales. For example, a smaller company may view a \$1M loss as Severe, and will need to adjust their loss scale accordingly.

A word about threat agents

Recall an earlier statement that not all threat agents are created equal. Within the context of analyzing loss magnitude, this means that the world has fewer mass murderers in it than jaywalkers – thankfully. Consequently, we can’t assume that all loss events are going to involve gross maliciousness. There are a couple of ways to deal with this as we perform our analyses:

- ! **Be very general** in how we define the threat community under analysis (e.g., “employees”), recognizing the probability of gross maliciousness within that population is low, or
- ! **Be more specific** in how we define the threat community under analysis (e.g., “disgruntled employees”), and recognize that this greater specificity provides a better understanding of motives (in this case, revenge), and an opportunity to be more precise in characterizing their likely actions (in this case, a much higher probability of significantly malicious actions). Note, however, that threat communities with higher propensity for gross maliciousness are typically much smaller, which affects the contact frequency.

The bottom line is that the better we understand the threat community under analysis, the more effective we can be at estimating loss.

Estimating loss

Other risk sciences often use very complex formulas for estimating loss probabilities. We'll keep our process relatively simple for now:

- 1) Evaluate the worst-case scenario
- 2) Evaluate the probable loss magnitude

Worst-case loss

Although we don't want to focus solely on worst-case loss, decision-makers need to know what a worst-case scenario might look like. It's also important to identify for them which key factors would drive a worst-case loss, as well as the likelihood of such an outcome.

Worst-case loss occurs when a specific set of factors converge. Because each of these factors has less than 100% probability of occurring, the probability of a worst-case outcome is the product of these probabilities. For example, if three factors were necessary in order to experience a worst-case outcome, and the individual probabilities for these factors were 25%, 50%, and 10%, then the worst-case probability would be: $.25 \times .5 \times .1 = .0125$ (just over 1%), assuming those factors vary independent of each other. Generally, the more factors required for worst-case, the lower the probability.

It isn't necessary to identify and estimate the probability for every conceivable factor. It's just important to recognize how few factors need to be in the mix in order to drive the probability of a worst-case outcome to a fraction of a percent.

We'll follow three steps to estimate worst-case magnitude:

- 1) Determine the threat action that would most likely result in worst-case
- 2) Estimate the magnitude for each loss form associated with that threat action
- 3) "Sum" the loss form magnitudes

For example, if our scenario involved sensitive information, a worst-case outcome might occur if the information was disclosed to the public. Referencing the loss factors we covered earlier, we evaluate the worst-case outcome by estimating the magnitude for each form of loss associated with disclosure.

Threat Actions	Loss Forms					
	Productivity	Response	Replacement	Fine/Judgments	Comp. Adv.	Reputation
Access						
Misuse						
Disclosure	M	H	VL	H	H	SV
Modification						
Deny Access						

Magnitude	Range Low End	Range High End
Severe (SV)	\$10,000,000	--
High (H)	\$1,000,000	\$9,999,999
Significant (Sg)	\$100,000	\$999,999
Moderate (M)	\$10,000	\$99,999
Low (L)	\$1,000	\$9,999
Very Low (VL)	\$0	\$999

We then “sum” the magnitudes to arrive at a worst-case outcome. In this case, our worst-case outcome is Severe. This may seem like a lot of work for an estimate we might be able to arrive at intuitively. There are, however, several advantages to this process:

- ! Improves consistency
- ! Helps to reduce individual bias
- ! Helps to ensure that we don’t overlook something important
- ! Provides a framework for explaining how we arrived at our conclusions
- ! Enables us to identify key factors and estimate the likelihood of a worst-case outcome

If our analysis identified four key factors that were required in order to experience a worst-case outcome, we could estimate the probability of those factors converging, and have a better understanding of worst-case likelihood.

Estimating probable loss magnitude (PLM)

There are three steps in our PLM evaluation process:

- 1) Identify the most likely threat community action(s)
- 2) Evaluate the probable loss magnitude for each loss form

3) Sum the magnitudes

When we evaluated worst-case, we selected the threat action that was likely to result in maximum loss. In this case, we select the action(s) the threat community is most likely to take. In some scenarios, the most likely threat action and worst-case action will be the same. In other scenarios they'll be different. We also may not always be able to pin our threat community down to a single most likely action. For example, if our threat community was defined as "Disgruntled Employees", we might reasonably expect the most likely actions to be Deny Access (e.g., destroy) in order to affect productivity, or Disclosure in an attempt to inflict legal/regulatory or reputational damage.

Referencing the loss factors shown below, we estimate the probable loss magnitude for each form of loss associated with the threat action(s) we've determined are most likely.

Threat Actions	Loss Forms					
	Productivity	Response	Replacement	Fine/Judgments	Comp. Adv.	Reputation
Access						
Misuse						
Disclosure						
Modification						
Deny Access						

Magnitude	Range Low End	Range High End
Severe (SV)	\$10,000,000	--
High (H)	\$1,000,000	\$9,999,999
Significant (Sg)	\$100,000	\$999,999
Moderate (M)	\$10,000	\$99,999
Low (L)	\$1,000	\$9,999
Very Low (VL)	\$0	\$999

NOTE: Including subject matter experts from the organization (e.g., legal counsel, business analysts, technology operations, etc.) can significantly improve the accuracy and credibility of your loss estimates.

Deriving and Articulating Risk

With everything we've covered so far, you might believe that deriving risk would be difficult. Recall that risk is simply the union of loss event frequency (LEF) and probable loss magnitude (PLM).




Figure 12

Once we've established those, we have risk. The real challenge has to do with articulating risk to our decision-makers.

Few of the decision-makers I'm familiar with feel as though a simple adjective-noun label (e.g., high-risk) provides enough information to make a well-informed decision. There's too much ambiguity. This is especially true if they believe that little or no depth of analysis exists underneath that label. These same executives generally wouldn't tolerate a simple "High Risk" label for product, market, or investment risk. Information risk should be no different.

Another challenge associated with simple qualitative labeling is that there's a tendency to equate "high-risk" with "unacceptable", and "low-risk" with "acceptable". The fact is, in some circumstances high-risk is entirely acceptable (e.g., in cases where the potential for reward outweighs the risk). In other situations, a relatively low-risk condition may be unacceptable, particularly if the exposure is systemic within an organization. Including more specific information regarding LEF and PLM can help to reduce the biases associated with qualitative risk labels.

The bottom line is that our risk articulation must meet the needs of the decision-makers. If they're comfortable with a simple qualitative label, so be it. The challenge then becomes making certain that they agree with the criteria for each level. The following matrix contains qualitative risk labels that may or may not represent the risk tolerances of your decision-makers.

		Risk				
PLM	Severe	H	H	C	C	C
	High	M	H	H	C	C
	Significant	M	M	H	H	C
	Moderate	L	M	M	H	H
	Low	L	L	M	M	M
	Very Low	L	L	M	M	M
		VL	L	M	H	VH
		LEF				

Key	Risk Level
C	Critical
H	High
M	Medium
L	Low

For example, some decision-makers may feel that a combination of Very Low LEF and High PLM represents High risk, as opposed to the Medium risk derived above. The lines that differentiate risk levels will vary from organization to organization.

Analyzing a Simple Scenario

We've covered a lot of ground, and it can be difficult to pull all of these concepts together until you've had an opportunity to use them. This section takes us through a simple risk scenario – providing an opportunity to kick the tires, so-to-speak.

The Scenario

A Human Resources (HR) executive within a large bank has his username and password written on a sticky-note stuck to his computer monitor. These authentication credentials allow him to log onto the network and access the HR applications he's entitled to use.

Before we get started, think to yourself how you'd rate the level of risk within this scenario based upon the assessments you've seen or done in the past.

The Analysis

The simplified process we'll use in this example is comprised of ten steps in four stages:

- ! Stage 1 – Identify scenario components
 - Identify the asset at risk
 - Identify the threat community under consideration
- ! Stage 2 – Evaluate Loss Event Frequency (LEF)
 - Estimate the probable Threat Event Frequency (TEF)
 - Estimate the Threat Capability (TCap)
 - Estimate Control strength (CS)
 - Derive Vulnerability (Vuln)
 - Derive Loss Event Frequency (LEF)
- ! Stage 3 – Evaluate Probable Loss Magnitude (PLM)
 - Estimate worst-case loss
 - Estimate probable loss
- ! Stage 4 – Derive and articulate Risk
 - Derive and articulate Risk

(Appendix A of this document contains a Basic Risk Assessment Guide that documents these steps.)

Stage 1 - Identify Scenario Components

Identify the Asset at Risk

The first question we have to answer is, “What asset is at risk?” Another way to think about this is to determine where value or liability exists. I’m typically asked when I present this scenario whether the credentials are the asset, or whether it’s the applications, systems, and information that the credentials provide access to. The short answer is “yes” – they’re all assets. In this case, however, we’ll focus on the credentials, recognizing that their value is inherited from the assets they’re intended to protect.

Identify the Threat Community

The second question we have to answer is, “Risk associated with what threat?” If we examine the nature of the organization (e.g., the industry it’s in, etc.), and the conditions surrounding the asset (e.g., an HR executive’s office), we can begin to parse the overall threat population into communities that might reasonably apply. How many threat communities we choose to analyze, and how we subdivide them is up to us. It’s probably not a good use of time to include every conceivable threat community in our analysis. For example, given this scenario, it probably wouldn’t be worthwhile to analyze the risk associated with nation-state intelligence services such as the French DGSE. Are we saying

that it's not possible for a nation-state spy to attack this bank through this exposure? No. But by considering the nature of the threat communities relative to the industry, organization, and asset, we can come to reasonable conclusions without falling victim to analysis paralysis or "lottery odds nit-picking".

Within this scenario, it seems reasonable to consider the risk associated with the following threat communities:

- ! The cleaning crew
- ! Other HR workers with regular access to the executive's office, and
- ! Visitors to the his office
- ! Guests
- ! Job applicants
- ! Technical support staff

With experience it becomes easier to determine which communities are worthwhile to include and exclude, and whether it makes sense to combine communities such as those that fall under 'Visitors'. For this example, let's focus on the cleaning crew.

Stage 2 – Evaluate Loss Event Frequency (LEF)

Estimate the probable Threat Event Frequency (TEF)

Many people demand reams of hard data before they're comfortable estimating attack frequency. Unfortunately, because we don't have much (if any) really useful or credible data for many scenarios, TEF is often ignored altogether. The minute we ignore this component of risk, however, we're no longer talking about risk. So, in the absence of hard data, what's left? One answer is to use a qualitative scale, such as Low, Medium, or High. And, while there's nothing inherently wrong with a qualitative approach in many circumstances, a quantitative approach provides better clarity and is more useful to most decision-makers – even if it's imprecise. For example, I may not have years of empirical data documenting how frequently cleaning crew employees abuse usernames and passwords on sticky-notes, but I can make a reasonable estimate within a set of ranges.

As discussed in the Factoring section, a TEF estimate would be based upon how frequently contact between this threat community (the cleaning crew) and the credentials occurs AND the probability that they would act against the credentials. If the cleaning crew comes by once per workday, contact reasonably occurs a couple of hundred times per year. The probability that they would act is driven by three primary factors:

- ! The value of the asset to them (based upon their motives – financial gain, revenge, etc.),
- ! How vulnerable the asset appears to be
- ! Versus the risk of being caught and suffering unacceptable consequences

Recognizing that cleaning crews are generally comprised of honest people, that an HR executive's credentials typically would not be viewed or recognized as especially valuable to them, and that the perceived risk associated with illicit use might be high, then it seems reasonable to estimate a **Low** TEF using the table below.

Rating	✓	Description
Very High (VH)		> 100 times per year
High (H)		Between 10 and 100 times per year
Moderate (M)		Between 1 and 10 times per year
Low (L)	✓	Between .1 and 1 times per year
Very Low (VL)		< .1 times per year (less than once every ten years)

Is it possible for a cleaning crew to have an employee with motive, sufficient computing experience to recognize the potential value of these credentials, and with a high enough risk tolerance to try their hand at illicit use? Absolutely! Does it happen? Undoubtedly. Might such a person be on the crew that cleans this office? Sure – it's possible. Nonetheless, the probable frequency is relatively low.

Estimate the Threat Capability (Tcap)

Tcap refers to the threat agent's skill (knowledge & experience) and resources (time & materials) that can be brought to bear against the asset. A different scenario might provide a better illustration of this component of the analysis – something like a web application with a SQL injection weakness – but scenarios like that don't lend themselves to an introductory document. In this case, all we're talking about is estimating the skill (in this case, reading ability) and resources (time) the average member of this threat community can use against a password written on a sticky note. It's reasonable to rate the cleaning crew Tcap as **Medium**, as compared to the overall threat population. Keep in mind that Tcap is always estimated relative to the scenario. If our scenario was different, and we were evaluating the cleaning crew's capability to execute a SQL injection attack, we'd probably rate them lower.

Rating	✓	Description
Very High (VH)		Top 2% when compared against the overall threat population
High (H)		Top 16% when compared against the overall threat population
Moderate (M)	✓	Average skill and resources (between bottom 16% and top 16%)
Low (L)		Bottom 16% when compared against the overall threat population
Very Low (VL)		Bottom 2% when compared against the overall threat population

Estimate the Control Strength (CS)

Control strength has to do with an asset's ability to resist compromise. In our scenario, because the credentials are in plain sight and in plain text, the CS is **Very Low**. If they were written down, but encrypted, the CS would be different – probably much higher.

Rating	✓	Description
Very High (VH)		Protects against all but the top 2% of an avg. threat population
High (H)		Protects against all but the top 16% of an avg. threat population
Moderate (M)		Protects against the average threat agent
Low (L)		Only protects against bottom 16% of an avg. threat population
Very Low (VL)	✓	Only protects against bottom 2% of an avg. threat population

The question sometimes comes up, “Aren't good hiring practices a control for internal assets?” and “Isn't the lock on the executive's door a control?” Absolutely, they are. But these controls factor into the frequency of contact, as opposed to how effective the controls are at the point of attack. We'll cover defense in depth in subsequent documentation.

Derive Vulnerability (Vuln)

Deriving vulnerability is easy once you've established your Tcap and CS. Recall from the Factoring section that vulnerability is the difference between the force that's likely to be applied, and the asset's ability to resist that force. Using the matrix below, simply find the Tcap along the left side of the matrix, and the CS along the bottom. Where they intersect determines the vulnerability.

		Vulnerability				
Tcap	VH	VH	VH	VH	H	M
	H	VH	VH	H	M	L
	M	VH	H	M	L	VL
	L	H	M	L	VL	VL
	VL	M	L	VL	VL	VL
		VL	L	M	H	VH
		Control Strength				

Derive Loss Event Frequency (LEF)

Similar to vulnerability, LEF is derived by intersecting TEF and Vulnerability within a matrix.

		Loss Event Frequency				
TEF	VH	M	H	VH	VH	VH
	H	L	M	H	H	H
	M	VL	L	M	M	M
	L	VL	VL	L	L	L
	VL	VL	VL	VL	VL	VL
		VL	L	M	H	VH
		Vulnerability				

In our scenario, given a TEF of Low and a Vuln of VH, the LEF is **Low**. Keep in mind that vulnerability is a percentage, which means that you can never be more than 100% vulnerable. Consequently, LEF will never be greater than TEF.

About now, you might start to feel uncomfortable with the idea of categorizing this scenario as a “low” anything – even if you agree with the logic behind the analysis. It runs counter to much of what we’ve been taught and have practiced. Questions regarding aggregate risk, and due diligence begin to arise – and for good reason. We’ll touch on this again later on, but rest assured that “low” LEF is not necessarily the same thing as “not a problem”.

Stage 3 – Evaluate Probable Loss Magnitude (PLM)

Using the previous seven steps, we’ve determined that the probability of a loss event in our scenario is Low (somewhere between .1 and 1 times per year). Now we’re faced with analyzing loss if an event does occur.

As mentioned earlier, the username and password credentials inherit the value and liability associated with the resources they provide access to. For an HR executive, we can reasonably expect these credentials to provide access to HR organizational information (org. charts, etc.), as well as employee personal and employment information (performance data, health and medical data, address, SSN, salary, etc.). In some organizations, depending upon where the HR executive exists in the corporate hierarchy, he/she might also have access to corporate strategy data. For our scenario, we’ll assume that this executive does not have access to key sensitive corporate strategies.

Estimate worst-case loss

Within this scenario, three potential threat actions stand out as having significant loss potential – misuse, disclosure, and destruction.

- ! **Misuse** – Employee records typically have information that can be used to execute identity theft, which introduces potential legal and reputational loss
- ! **Disclosure** – Employee records often have sensitive personal information related to medical or performance issues, which introduces legal and reputational exposure
- ! **Deny access (destruction)** – Employee records are a necessary part of operating any business. Consequently, their destruction can introduce some degree of lost productivity.

In some cases it's necessary to evaluate the loss associated with more than one threat action in order to decide which one has the most significant loss potential. For this exercise, we'll select disclosure as our worst-case threat action.

Our next step is to estimate the worst-case loss magnitude for each loss form.

Threat Actions	Loss Forms					
	Productivity	Response	Replacement	Fine/Judgments	Comp. Adv.	Reputation
Access						
Misuse						
Disclosure	H	H	--	SV	H	SV
Modification						
Deny Access						

Magnitude	Range Low End	Range High End
Severe (SV)	\$10,000,000	--
High (H)	\$1,000,000	\$9,999,999
Significant (Sg)	\$100,000	\$999,999
Moderate (M)	\$10,000	\$99,999
Low (L)	\$1,000	\$9,999
Very Low (VL)	\$0	\$999

Note that we didn't estimate loss magnitude for Replacement. Any time you're evaluating loss and one or more of the forms has a loss magnitude of Severe (Sv), it's not worthwhile to give much thought to loss forms having a much lower, or no, loss magnitude. In this case, Replacement doesn't apply because the assets aren't being destroyed.

Our estimates are based on the following rationale:

- ! Productivity
 - It's conceivable that productivity losses could be High as employee attention is diverted to this event
- ! Response
 - Legal expenses associated with inside and outside legal counsel could be High, particularly if class action lawsuits were filed
- ! Fines/Judgments
 - If the disclosed information included details regarding psychological illness or other sensitive health issues, then legal judgments in behalf of affected employees could be Severe, particularly if a large number of employees were affected
 - If the information included evidence of criminal activity or incompetence on the part of management, then legal and regulatory fines and sanctions could be Severe
- ! Competitive advantage
 - If the disclosed information provided evidence of incompetence or criminal activity, competitors could, in theory, leverage that to gain advantage. For the most part, however, we can expect competitors to simply sit back and rake in any disaffected customers (falls under reputational loss)
- ! Reputation
 - If the information was sensitive enough, due diligence was seriously absent, legal actions were large enough, and media response was negative and pervasive, then reputational loss associated with customer flight and stock value could be Severe.

** Magnitudes will vary based on the size of the organization.*

We aren't going to document all of our rationale in most risk analyses. Most of the time we internalize all but the most significant factors. Nonetheless, having a deeper understanding of what these factors are and how they work increases the quality of our analyses.

Note that the rationale above is based on what could happen. This highlights the fact that worst-case analyses tend to be based on possibilities rather than probabilities. In order to make this worst-case information meaningful, we need to have some idea of how probable a worst-case outcome is.

A large number of factors affect the likelihood of a worst-case outcome. In this scenario, we selected disclosure as our worst-case threat action, yet we haven't considered the likelihood that a threat agent from this threat community would intentionally disclose the information. Other actions might be far more likely. Accidental disclosure might result, of course, if the threat agent performed identity theft, was caught, and the information was traced back to this organization and this event. A series of 'ifs' – each with less than a 100% probability. Furthermore, even if disclosure occurred, the organization has an opportunity to mitigate loss magnitude through its response. Does it go out of its way to rectify the situation? Does it have an effective public relations capability and a good relationship with the media? Each of these factors reduce the probability of a worst-case outcome.

In most cases it isn't worthwhile to spend too much time and effort evaluating the probability of a worst-case outcome. Spend enough time to get a sense for what the key factors are, and roughly where on the continuum worst-case outcome falls between almost certain and almost impossible.

For our scenario, we'll determine that worst-case magnitude is severe (tens of millions of dollars), but with a very low probability of occurring.

Estimate probable loss magnitude (PLM)

The first step in estimating PLM is to determine which threat action is most likely. Remember; actions are driven by motive, and the most common motive for illicit action is financial gain. Given this threat community, the type of asset (personal information), and the available threat actions, it's reasonable to select Misuse as the most likely action – e.g., for identity theft.

Our next step is to estimate the most likely loss magnitude resulting from Misuse for each loss form.

Threat Actions	Loss Forms					
	Productivity	Response	Replacement	Fine/Judgments	Comp. Adv.	Reputation
Access						
Misuse	M	M	VL	VL	VL	VL
Disclosure						
Modification						
Deny Access						

Magnitude	Range Low End	Range High End
Severe (SV)	\$10,000,000	--
High (H)	\$1,000,000	\$9,999,999
Significant (Sg)	\$100,000	\$999,999
Moderate (M)	\$10,000	\$99,999
Low (L)	\$1,000	\$9,999
Very Low (VL)	\$0	\$999

Our rationale for these estimates include:

- ! The impact to productivity will be Moderate as employees react to the event
- ! The cost of responding to the event will include investigation, some amount of time from internal legal counsel, and providing restitution to any affected employees
- ! Replacement expenses simply entail the cost of changing the executive's password
- ! No legal or regulatory action occurs because the incident isn't taken to court or reported to the regulators
- ! No competitive advantage loss occurs due to the relatively inconsequential nature of the event
- ! No material reputational damage occurs because it was an internal event, no customers were affected, and the organization had a security program in place that included policies and education

A few key assumptions also played a role in our estimates. We assumed:

- ! The organization became aware of the incident. It's entirely possible for this kind of event to go undetected. Until detected, there is no material loss to the organization.
- ! Relatively few employees actually experienced identity theft.
- ! The organization responded effectively to the event.

Stage 4 – Derive and Articulate Risk

Derive and Articulate Risk

We've already done the hard part, as risk is simply derived from LEF and PLM. The question is whether to articulate risk qualitatively using a matrix like the one below, or articulate risk as LEF, PLM, and worst-case. For this exercise, we'll do both.

Assuming that the matrix below has been 'approved' by the leadership of our fictional bank, we can report that risk associated with this threat community is Medium based upon a low LEF (between 1 and .1 times per year) and a moderate PLM (between \$10K and \$100K). Furthermore, we can communicate to our decision-makers that worst-case loss could be severe, but that the probability of a worst-case outcome is very low.

		Risk				
PLM	Severe	H	H	C	C	C
	High	M	H	H	C	C
	Significant	M	M	H	H	C
	Moderate	L	M	M	H	H
	Low	L	L	M	M	M
	Very Low	L	L	M	M	M
		VL	L	M	H	VH
		LEF				

Key	Risk Level
C	Critical
H	High
M	Medium
L	Low

In a real analysis, it's likely that we would evaluate and report on more than one threat community.

A word of caution: Although the risk associated with any single exposure may be relatively low, that same exposure existing in many instances across an organization may represent a higher aggregate risk. Under some conditions, the aggregate risk may increase geometrically as opposed to linearly. Furthermore, low risk issues, of the wrong types and in the wrong combinations, may create an environment where a single event can cascade into a catastrophic outcome – an avalanche effect. It's important to keep these considerations in mind when evaluating risk and communicating the outcome to decision-makers. Subsequent FAIR documentation and training will cover these issues in detail.

Conclusions and Next Steps

Conclusions

Our profession has recognized all along that perfect security isn't possible, nor would it be practical if it were possible. Our fundamental purpose as professionals is to help our employers manage the frequency and magnitude of loss. Unfortunately, the methods and concepts many of us have followed for years don't reflect the true nature of risk, and have limited our ability to be effective. We haven't been able to credibly answer some very basic questions:

- ! How much risk management is enough?
- ! How much risk do we have?
- ! How much less risk will we have if we employ solution X, Y, or Z?


Each of these questions implies an ability to measure risk. Yet without a solid understanding of fundamental risk concepts and factors, we can't begin to credibly measure it. FAIR seeks to provide the necessary foundation through its taxonomy, definitions, and analysis methods.

As I've incubated and applied these concepts and processes to real-world risk scenarios where I work, and as I've begun to train others within my organization, the results have been significant:

- ! Much more consistent, higher quality analyses
- ! A greater feeling of confidence by those performing the analyses
- ! A greater feeling of confidence by decision-makers
- ! A significantly improved ability to cost-effectively manage risk

Of those who have been introduced to FAIR and have had a chance to see it used, most are very enthusiastic in their support. It's natural, though, for people to accept change at different speeds. Some of us hold our beliefs very firmly, and it can be difficult and uncomfortable to adopt a new approach. Ultimately, not everyone is going to agree with the principles or methods that underlie FAIR. A few have called it nonsense. Others appear to feel threatened by it. Their concerns tend to revolve around one or more of the following issues:

- ! **The absence of hard data.** There's no question that an abundance of good data would be useful. Unfortunately, that's not our current reality. Consequently, we need to find another way to approach the problem, and FAIR is one solution.
- ! **The lack of precision.** Here again, precision is nice when it's achievable, but it's not realistic within this problem space. Reality is just too complex. Consider the following illustrations:



The target on the left has a relatively precise shot pattern, but the placement isn't accurate. The target on the right has a less precise shot pattern, but the placement reflects far better accuracy. Many of the assessment methods and best practices used today provide a relatively high degree of precision, but only address control state. Unfortunately, control state often doesn't accurately reflect risk. **FAIR represents an attempt to gain far better accuracy, while recognizing that the fundamental nature of the problem doesn't allow for a high degree of precision.** My experience has been that decision makers strongly prefer accuracy.

- ! **It takes some of the mystery out of the profession.** The fact is, there are those who prefer to be artists – in some cases because an artist can never be judged as “wrong.”
- ! **FAIR analysis appears to be hard work.** The good news is that it gets easier with practice. After awhile, our quick, “intuitively guided” risk judgments become much higher quality because of our deeper understanding. We're better calibrated. It's also worth noting that even simple prototype FAIR software applications make complex analyses significantly easier.
- ! **FAIR appears complicated.** There's no question that most of us like simple solutions when we can find them. In fact, simple solutions are more effective than complex solutions in many cases. Fortunately, we can choose to use the framework at whatever level of abstraction suits our needs. The advantage comes from knowing more about factors that exist at lower levels of abstraction, which enables us to make better judgments at higher levels of abstraction.
- ! **Some people just don't like change** – particularly change as profound as this represents.

It isn't surprising that some people react negatively, because FAIR represents a disruptive influence within our profession. My only request of those who offer criticism is that they also offer rational reasons and alternatives. In fact, I encourage hard questions and constructive criticism because:

- ! Weaknesses in the framework can be identified and corrected, or
- ! The framework may provide an answer to the question or criticism, which strengthens its credibility

Where to go from here

This document barely scratches the surface of development and research spanning four years. More comprehensive documentation, prototype tools, and risk analysis training materials are being developed so that the framework can be applied against the complex real-world issues we face every day. Future documentation will cover the following topics:

- ! A much deeper dive into controls, threat communities, and loss
- ! Data capture and analysis
- ! Complex scenario modeling
- ! Broad-spectrum threat analysis
- ! Fragility and instability concepts
- ! Aggregate risk
- ! Error, failure, and acts of God
- ! Evaluating risk at the organizational level
- ! Integrating FAIR concepts into an organizational risk management program
- ! Using FAIR concepts to evaluate other types of risk (e.g., market risk, investment risk, legal risk, etc.)

Your feedback, questions, and insights are most welcome, and I will respond in as timely a manner as workload permits. Please submit e-mails to:

jonesj1@riskmanagementinsight.com

Include “FAIR” in the subject line.

“...and the end of all of our exploring will be to arrive where we started and know the place for the first time.”

- T.S. Eliot

Appendix A: Basic Risk Assessment Guide

NOTE: Before using this assessment guide...

Using this guide effectively requires a solid understanding of FAIR concepts

- ! As with any high-level analysis method, results can depend upon variables that may not be accounted for at this level of abstraction
- ! The loss magnitude scale described in this section is adjusted for a specific organizational size and risk capacity. Labels used in the scale (e.g., “Severe”, “Low”, etc.) may need to be adjusted when analyzing organizations of different sizes
- ! This process is a simplified, introductory version that may not be appropriate for some analyses

Basic FAIR analysis is comprised of ten steps in four stages:

Stage 1 – Identify scenario components

1. Identify the asset at risk
2. Identify the threat community under consideration

Stage 2 – Evaluate Loss Event Frequency (LEF)

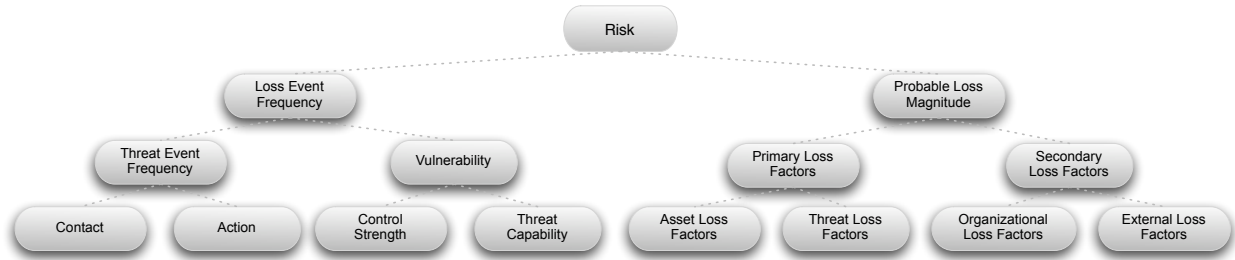
3. Estimate the probable Threat Event Frequency (TEF)
4. Estimate the Threat Capability (TCap)
5. Estimate Control strength (CS)
6. Derive Vulnerability (Vuln)
7. Derive Loss Event Frequency (LEF)

Stage 3 – Evaluate Probable Loss Magnitude (PLM)

8. Estimate worst-case loss
9. Estimate probable loss

Stage 4 – Derive and articulate Risk

10. Derive and articulate Risk



Stage 1 – Identify Scenario Components

Step 1 – Identify the Asset(s) at risk

In order to estimate the control and value characteristics within a risk analysis, the analyst must first identify the asset (object) under evaluation. If a multilevel analysis is being performed, the analyst will need to identify and evaluate the primary asset (object) at risk and all meta-objects that exist between the primary asset and the threat community. This guide is intended for use in simple, single level risk analysis, and does not describe the additional steps required for a multilevel analysis.

Asset(s) at risk: _____

Step 2 – Identify the Threat Community

In order to estimate Threat Event Frequency (TEF) and Threat Capability (TCap), a specific threat community must first be identified. At minimum, when evaluating the risk associated with malicious acts, the analyst has to decide whether the threat community is human or malware, and internal or external. In most circumstances, it's appropriate to define the threat community more specifically – e.g., network engineers, cleaning crew, etc., and characterize the expected nature of the community. This document does not include guidance in how to perform broad-spectrum (i.e., multi-threat community) analyses.

Threat community: _____

Characterization

Characterization	
------------------	--

Stage 2 – Evaluate Loss Event Frequency

Step 3 – Threat Event Frequency (TEF)

The probable frequency, within a given timeframe, that a threat agent will act against an asset

Contributing factors: Contact Frequency, Probability of Action

Rating	✓	Description
Very High (VH)		> 100 times per year
High (H)		Between 10 and 100 times per year
Moderate (M)		Between 1 and 10 times per year
Low (L)		Between .1 and 1 times per year
Very Low (VL)		< .1 times per year (less than once every ten years)

Rationale	
------------------	--

Step 4 – Threat Capability (Tcap)

The probable level of force that a threat agent is capable of applying against an asset

Contributing factors: Skill, Resources

Rating	✓	Description
Very High (VH)		Top 2% when compared against the overall threat population
High (H)		Top 16% when compared against the overall threat population
Moderate (M)		Average skill and resources (between bottom 16% and top 16%)
Low (L)		Bottom 16% when compared against the overall threat population
Very Low (VL)		Bottom 2% when compared against the overall threat population

Rationale	
------------------	--

Step 5 – Control strength (CS)

The expected effectiveness of controls, over a given timeframe, as measured against a baseline level of force

Contributing factors: Strength, Assurance

Rating	✓	Description
Very High (VH)		Protects against all but the top 2% of an avg. threat population
High (H)		Protects against all but the top 16% of an avg. threat population
Moderate (M)		Protects against the average threat agent
Low (L)		Only protects against bottom 16% of an avg. threat population
Very Low (VL)		Only protects against bottom 2% of an avg. threat population

Rationale

Step 6 – Vulnerability (Vuln)

The probability that an asset will be unable to resist the actions of a threat agent

Tcap (from step 4): _____

CS (from step 5): _____

		Vulnerability				
Tcap	VH	VH	VH	VH	H	M
	H	VH	VH	H	M	L
	M	VH	H	M	L	VL
	L	H	M	L	VL	VL
	VL	M	L	VL	VL	VL
		VL	L	M	H	VH
Control Strength						

Vuln (from matrix above): _____

Step 7 – Loss Event Frequency (LEF)

The probable frequency, within a given timeframe, that a threat agent will inflict harm upon an asset

TEF (from step 3): _____

Vuln (from step 6): _____

		Loss Event Frequency				
TEF	VH	M	H	VH	VH	VH
	H	L	M	H	H	H
	M	VL	L	M	M	M
	L	VL	VL	L	L	L
	VL	VL	VL	VL	VL	VL
		VL	L	M	H	VH
		Vulnerability				

LEF (from matrix above): _____

Stage 3 – Evaluate Probable Loss Magnitude

Step 8 – Estimate worst-case loss

Estimate worst-case magnitude using the following three steps:

- ! Determine the threat action that would most likely result in a worst-case outcome
- ! Estimate the magnitude for each loss form associated with that threat action
- ! “Sum” the loss form magnitudes

Threat Actions	Loss Forms					
	Productivity	Response	Replacement	Fine/Judgments	Comp. Adv.	Reputation
Access						
Misuse						
Disclosure						
Modification						
Deny Access						

Magnitude	Range Low End	Range High End
Severe (SV)	\$10,000,000	--
High (H)	\$1,000,000	\$9,999,999
Significant (Sg)	\$100,000	\$999,999
Moderate (M)	\$10,000	\$99,999
Low (L)	\$1,000	\$9,999
Very Low (VL)	\$0	\$999

Step 9 – Estimate probable loss

Estimate probable loss magnitude using the following three steps:

- ! Identify the most likely threat community action(s)
- ! Evaluate the probable loss magnitude for each loss form
- ! “Sum” the magnitudes

Threat Actions	Loss Forms					
	Productivity	Response	Replacement	Fine/Judgments	Comp. Adv.	Reputation
Access						
Misuse						
Disclosure						
Modification						
Deny Access						

Magnitude	Range Low End	Range High End
Severe (SV)	\$10,000,000	--
High (H)	\$1,000,000	\$9,999,999
Significant (Sg)	\$100,000	\$999,999
Moderate (M)	\$10,000	\$99,999
Low (L)	\$1,000	\$9,999
Very Low (VL)	\$0	\$999

Stage 4 – Derive and Articulate Risk

Step 10 – Derive and Articulate Risk

The probable frequency and probable magnitude of future loss

Well-articulated risk analyses provide decision-makers with at least two key pieces of information:

- ! The estimated loss event frequency (LEF), and
- ! The estimated probable loss magnitude (PLM)

This information can be conveyed through text, charts, or both. In most circumstances, it's advisable to also provide the estimated high-end loss potential so that the decision-maker is aware of what the worst-case scenario might look like. Depending upon the scenario, additional specific information may be warranted if, for example:

- ! Significant due diligence exposure exists
- ! Significant reputation, legal, or regulatory considerations exist

		Risk				
PLM	Severe	H	H	C	C	C
	High	M	H	H	C	C
	Significant	M	M	H	H	C
	Moderate	L	M	M	H	H
	Low	L	L	M	M	M
	Very Low	L	L	M	M	M
		VL	L	M	H	VH
		LEF				

LEF (from step 7): _____

PLM (from step 9): _____

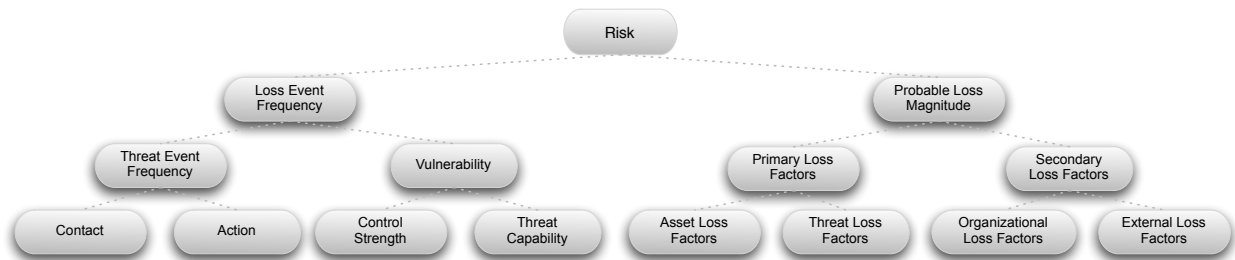
WCLM (from step 8): _____

Key	Risk Level
C	Critical
H	High
M	Medium
L	Low

Appendix B: Glossary

Term	Definition
Action	An act taken against an asset by a threat agent. Requires first that contact occur between the asset and threat agent.
Broad spectrum risk analysis	Any analysis that accounts for the risk from multiple threat communities against a single asset
Contact	Occurs when a threat agent establishes a physical or virtual (e.g., network) connection to an asset
Control strength (CS)	The strength of a control as compared to a standard measure of force
Loss event	Occurs when a threat agent's action (threat event) is successful in negatively affecting an asset
Loss event frequency (LEF)	The probable frequency, within a given timeframe, that a threat agent will inflict harm upon an asset
Multilevel risk analysis	Any analysis that accounts for the risk from a single threat community against a layered set of assets (e.g., defense in depth)
Probable loss magnitude (PLM)	The probable magnitude of loss resulting from a loss event
Risk	The probable frequency and probable magnitude of future loss
Threat agent	Any agent (e.g., object, substance, human, etc.) that is capable of acting against an asset in a manner that can result in harm
Threat capability (Tcap)	The probable level of force that a threat agent is capable of applying against an asset
Threat community	A subset of the overall threat agent population that shares key characteristics
Threat event	Occurs when a threat agent acts against an asset
Threat event frequency (TEF)	The probable frequency, within a given timeframe, that a threat agent will act against an asset
Vulnerability	The probability that an asset will be unable to resist the actions of a threat agent

Appendix C: Factoring Diagram



About the Author

Jack has been employed in information technology since 1983, and has specialized in information security management, consulting, and assessment since 1991. His experience spans the military, government intelligence, consulting, as well as the commercial insurance and finance industries. In 2006 Jack was honored to receive the ISSA Excellence in the Field of Information Security Practices award, and in 2007 he was selected as a finalist for the Information Security Executive of the Year, Central United States. He has contributed to various publications, and is a sought-after speaker for national conferences.