Technical Standard

Risk Taxonomy

Technical Standard

Risk Taxonomy

ISBN: 1-931624-77-1

Document Number: C081

Published by The Open Group, January 2009.

Comments relating to the material contained in this document may be submitted to:

The Open Group
Thames Tower
37-45 Station Road
Reading
Berkshire, RG1 1LX
United Kingdom

or by electronic mail to:

ogspecs@opengroup.org

Contents

Preface

The Open Group

The Open Group is a vendor-neutral and technology-neutral consortium, whose vision of Boundaryless Information Flow™ will enable access to integrated information within and between enterprises based on open standards and global interoperability. The Open Group works with customers, suppliers, consortia, and other standards bodies. Its role is to capture, understand, and address current and emerging requirements, establish policies, and share best practices; to facilitate interoperability, develop consensus, and evolve and integrate specifications and Open Source technologies; to offer a comprehensive set of services to enhance the operational efficiency of consortia; and to operate the industry's premier certification service, including UNIX® certification.

Further information on The Open Group is available at www.opengroup.org.

The Open Group has over 15 years' experience in developing and operating certification programs and has extensive experience developing and facilitating industry adoption of test suites used to validate conformance to an open standard or specification.

More information is available at www.opengroup.org/certification.

The Open Group publishes a wide range of technical documentation, the main part of which is focused on development of Technical and Product Standards and Guides, but which also includes white papers, technical studies, branding and testing documentation, and business titles. Full details and a catalog are available at www.opengroup.org/bookstore.

As with all *live* documents, Technical Standards and Specifications require revision to align with new developments and associated international standards. To distinguish between revised specifications which are fully backwards-compatible and those which are not:

- A new *Version* indicates there is no change to the definitive information contained in the previous publication of that title, but additions/extensions are included. As such, it *replaces* the previous publication.
- A new *Issue* indicates there is substantive change to the definitive information contained in the previous publication of that title, and there may also be additions/extensions. As such, both previous and new documents are maintained as current publications.

Readers should note that updates – in the form of Corrigenda – may apply to any publication. This information is published at www.opengroup.org/corrigenda.

This Document

This document is the Technical Standard for Risk Taxonomy. It has been developed and approved by The Open Group.

This document provides a standard definition and taxonomy for information security risk, as well as information regarding how to use the taxonomy.

The intended audience for this document includes anyone who needs to understand and/or analyze a risk condition. This includes, but is not limited to:

- Information security and risk management professionals
- Auditors and regulators
- Technology professionals
- Management

Note that this taxonomy is not limited to application in the information security space. It can, in fact, be applied to any risk scenario. This agnostic characteristic enables the taxonomy to be used as a foundation for normalizing the results of risk analyses across varied risk domains.

This Risk Taxonomy Technical Standard is the first in an initial set of three Open Group publications addressing Risk Management. Following publications will be:

- **The Open Group Technical Guide: Requirements for Risk Assessment Methodologies** identifies and describes the key characteristics that make up any effective risk assessment methodology, thus providing a common set of criteria for evaluating any given risk assessment methodology against a clearly defined common set of essential requirements. In this way, it explains what features to look for when evaluating the capabilities of any given methodology, and the value those features represent.
- **The Open Group Technical Standard: Risk Assessment Methodology & Cookbook** describes in detail how to apply the FAIR (Factor Analysis for Information Risk) methodology to a selected risk management framework, in the form of an application paper. FAIR is complementary to other methodologies like COSO, ITIL, ISO/IEC 27002:2005, COBIT, OCTAVE, etc. – it provides the engine that can be used in other risk models. The Cookbook part of this document enables risk technology practitioners to follow by example how to create their own application to apply FAIR to other frameworks of their choice.

Trademarks

Acknowledgements

The Open Group gratefully acknowledges the contribution of Alex Hutton and Jack Jones, Risk Management Insight, for their valued original work and their continued support in guiding the Security Forum members through our process to develop this Risk Taxonomy Technical Standard.

The Open Group also acknowledges the members of the Security Forum who have contributed to the development of this Technical Standard.

Referenced Documents

The following documents are referenced in this Technical Standard:

- An Introduction to Factor Analysis of Information Risk (FAIR), Risk Management Insight LLC, November 2006; refer to: www.riskmanagementinsight.com.
- Methods for the Identification of Emerging and Future Risks, European Network and Information Security Agency (ENISA), November 2007; refer to www.enisa.europa.eu/doc/pdf/deliverables/EFR_Methods_Identification_200804.pdf.
- Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE), US-CERT; refer to www.cert.org/octave.
- A Taxonomy of Computer Program Security Flaws, with Examples, Naval Research Laboratory, September 1994; refer to: http://chacs.nrl.navy.mil/publications.

1 Introduction

1.1 Scope

This Technical Standard provides a taxonomy describing the factors that drive risk – their definitions and relationships.

This Technical Standard is not a reference or tutorial on how to assess or analyze risk, as there are many such references already available. This Technical Standard also does not cover those elements of risk management that pertain to strategic and tactical risk decisions and execution.

In the overall context of risk management, it is important to appreciate that our business objective in performing risk assessments is to identify and estimate levels of exposure to the likelihood of loss, so that business managers can make informed business decisions on how to manage those risks of loss – either by accepting each risk, or by mitigating it – through investing in appropriate internal protective measures judged sufficient to lower the potential loss to an acceptable level, or by investing in external indemnity. Critical to enabling good business decision-making therefore is to use risk assessment methods which give objective, meaningful, consistent results.

Fundamental to risk assessments is a sound approach:

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

The problem here is that a variety of definitions do exist, but the risk management community has not yet adopted a consistent definition for even the most fundamental terms in its vocabulary; e.g., threat, vulnerability, even risk itself. Without a sound common understanding of what risk is, what the factors are that drive risk, and a standard use of the terms we use to describe it, we can't be effective in delivering meaningful, comparable risk assessment results. This Risk Taxonomy provides the necessary foundation vocabulary, based on a fundamental analysis of what risk is, and then shows how to apply it to produce the objective, meaningful, and consistent results that business managers need.

1.2 Purpose/Objective

The purpose and objective of this Technical Standard is to provide a single logical and rational taxonomical framework for anyone who needs to understand and/or analyze information security risk. It can and should be used to:

- Educate information security, risk, and audit professionals
- Establish a common language for the information security and risk management profession

- Introduce rigor and consistency into analysis, which sets the stage for more effective risk modeling
- Explain the basis for risk analysis conclusions
- Strengthen existing risk assessment and analysis methods
- Create new risk assessment and analysis methods
- Evaluate the efficacy of risk assessment and analysis methods
- Establish metric standards and data sources

1.3 Context

Although the terms “risk” and “risk management” mean different things to different people, this Technical Standard is intended to be applied toward the problem of managing the frequency and magnitude of loss that arises from a threat (whether human, animal, or natural event). In other words, managing “how often bad things happen, and how bad they are when they occur”.

Although the concepts and taxonomy within this Technical Standard were not developed with the intention of being applied towards other risk types, experience has demonstrated that they can be effectively applied to other risk types. For example, they have been successfully applied in managing the likelihood and consequence of adverse events associated with project management or finance, in legal risk, and by statistical consultants in cases where probable impact is a concern (e.g., introducing a non-native species into an ecosystem).

1.4 The Risk Language Gap

Over time, the ways we manage risk have evolved to keep up with ways we conduct business. There is a very long history here, pre-dating the use of IT in business. As the scope, scale, and value of business operations have evolved, our specializations to manage the risk have similarly evolved, but in doing so each specialization has developed its own view of risk and how to describe its components. This has resulted in a significant language gap between the different specializations, all of whom are stakeholders in managing risk.

This gap is particularly evident between business managers and their IT risk/security specialists/analysts. For example, business managers talk about “impact” of loss not in terms of how many servers or operational IT systems will cease to provide normal service, but rather what will be the impact of losing these normal services on the business’s capacity to continue to trade normally, measured in terms of $-value; or will the impact be a failure to satisfy applicable regulatory requirements which could force them to limit or even cease trading and perhaps become liable to heavy legal penalties.

So, a business manager tends to think of a “threat” as something which could result in a loss which the business cannot absorb without seriously damaging its trading position. Compare this with our Risk Taxonomy definitions for “threat” and “vulnerability”:

Threat	Anything that is capable of acting in a manner resulting in harm to an asset and/or organization; for example, acts of God (weather, geological events, etc.); malicious actors; errors; failures.
Vulnerability	The probability that threat capability exceeds the ability to resist the threat.

Similar language gaps exist between other stakeholders in management of risk. Politicians and lawyers are particularly influential stakeholders. They are in the powerful position of shaping national and international policy (e.g., OECD, European Commission) which in turn influences national governments to pass laws and regulatory regimes on business practices that become effective one to three years down the line.

This Risk Taxonomy is an essential step towards enabling all stakeholders in risk management to use key risk management terms – especially Control, Asset, Threat, and Vulnerability – with precise meanings so we can bridge the language gap between IT specialists, business managers, lawyers, politicians, and other professionals, in all sectors of industry and commerce and the critical infrastructure, whose responsibilities bear on managing risk.

1.5 Using FAIR with Other Risk Assessment Frameworks

As The Open Group seeks to further its risk management framework based on FAIR, it is important to understand what the strengths of a FAIR approach are, and how they complement the work of other standards bodies. This section explains the outputs of a FAIR analysis and how these outputs are valuable in augmenting other risk assessment frameworks.

A valuable starting point here is the work published by the European Network and Information Security Agency (ENISA) in its November 2007 paper: Methods for the identification of Emerging and Future Risks. This ENISA document described how 18 various risk assessment frameworks addressed the criteria that the agency thought were important in assessing risk, and graded them on a numerical scale. In reviewing ENISA's criteria, the rating they assigned to each one, and the other risk assessment frameworks they reviewed, it became obvious that FAIR is not in direct competition with the other risk assessment frameworks, but actually is complementary to many of them.

1.5.1 The Ability of a FAIR-Based Approach to Complement Other Standards

FAIR, as a taxonomy of the factors that contribute to risk and how they affect each other, is primarily concerned with establishing accurate probabilities for the frequency and magnitude of loss events. It is not, *per se*, a "cookbook" that describes how to perform an enterprise (or individual) risk assessment. For example, FAIR documentation isn't so much concerned about the where and how you should get prior information for use in the assessment, as much as explaining how to describe the value of that information and how it contributes to creating risk.

So many risk assessment methodologies don't focus or concern themselves with how to establish consistent, defensible belief statements about risk – they simply give you steps they believe an organization should perform in order to have information for use in the creation of risk statements. As such, FAIR can be utilized within the context of many of these standards without significant modifications to FAIR or the other methodology.

1.5.2 An Example: Using FAIR with OCTAVE

One good example might be using FAIR to augment an OCTAVE (Operationally Critical Threat, Asset, and Vulnerability Evaluation) assessment. OCTAVE is a risk assessment methodology developed and sold by US-CERT (refer to www.cert.org/octave). In Version 2 of the OCTAVE criteria, the document authors mention at least three times that: "Using probability … is optional". Section 3.2 of OCTAVE then directs assessors to establish their own criteria and context for developing values (high, medium, low) for "impact" and "likelihood". Unfortunately, OCTAVE gives no structured means to determine why likelihood might be "high" or why impact might be "low". OCTAVE simply states:

"It is important to establish criteria (for the qualitative expressions) that are meaningful to the organization."

Practitioners who want a means to develop "meaningful" risk statements using FAIR would simply use the FAIR taxonomy and framework to build consistent and defensible risk statements. This could be accomplished by augmenting Section 3 of the OCTAVE criteria with the relevant parts of the FAIR basic risk assessment methodology (see Chapter 6) which describes how FAIR's basic risk assessment methodology comprises ten steps in four stages. In this example, the risk criteria in Section 3.2 of the OCTAVE criteria would be strengthened by using the appropriate steps in the FAIR basic risk assessment methodology, and the statement of risk required by Section 3.3 of the OCTAVE criteria would similarly be able to use the appropriate step in the FAIR methodology.

1.5.3 Conclusion

Just by glancing through the relevant parts of the ENISA document, an experienced FAIR practitioner can identify several other methodologies that FAIR complements (NIST 800-30, ISO/IEC 27002:2005, COBIT, ITIL, for example). FAIR also complements risk assessment frameworks not included in the ENISA document (for example, COSO; refer to www.coso.org/-ERM.htm). In fact, there are no commonly used methodologies for performing or communicating risk that would be antagonistic to the use of FAIR.

As a standards body, The Open Group aims to evangelize the use of FAIR within the context of these risk assessment or management frameworks. In doing so, The Open Group becomes not just a group offering yet another risk assessment framework, but a standards body which solves the difficult problem of developing consistent, defensible statements concerning risk.

2 Business Case

Risk management is fundamentally about making decisions – decisions about which risk issues are most critical (prioritization), which risk issues are not worth worrying about (risk acceptance), and how much to spend on the risk issues that need to be dealt with (budgeting). In order to be consistently effective in making these decisions, we need to be able to compare the issues themselves, as well as the options and solutions that are available. In order to compare, we need to measure, and measurement is predicated upon a solid definition of the things to be measured. Figure 1 shows these chained dependencies.

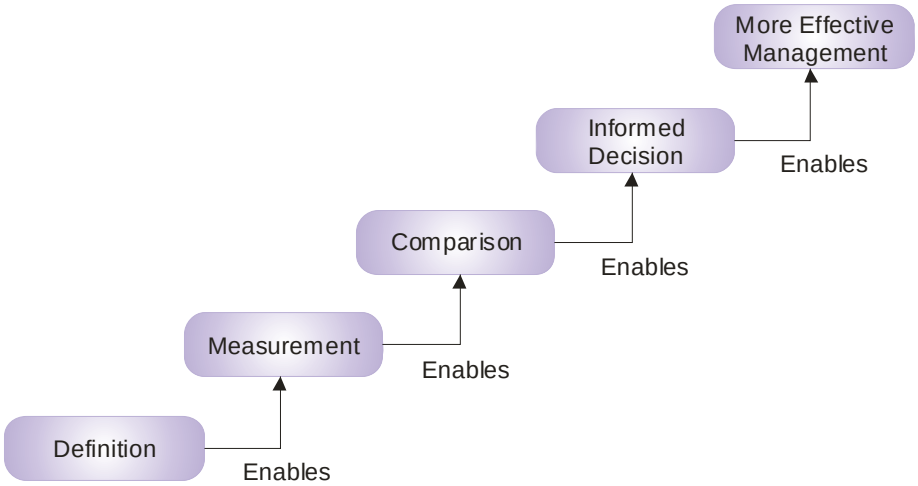

Figure 1

To date, the information security profession has been hamstrung by several challenges, not the least of which is inconsistent nomenclature. For example, in some references, software flaws/faults that could be exploited will be called a "threat", while in other references these same software faults will be referred to as a "risk", and yet other references will refer to them as "vulnerabilities". Besides the confusion that can result, this inconsistency makes it difficult if not impossible to normalize data and develop good metrics.

A related challenge stems from mathematical equations for risk that are either incomplete or illogical. For example, one commonly cited equation for risk states that:

*Risk = (Threat * Vulnerability) / Controls*

Amongst other problems, this equation doesn't tell us whether *Threat* means the level of force being applied or the frequency with which threat events occur. Furthermore, impact (magnitude of loss) is left out of the equation altogether. As we will touch on shortly, organization management cares very deeply about the question of loss magnitude, and so any risk equation that ignores impact is going to be meaningless to the very people who need to use risk analyses to make risk decisions.

These issues have been a major contributor to why the information security profession has consistently been challenged to find and maintain "a seat at the table" with the other organizational functions (e.g., finance, marketing, etc.). Furthermore, while few people are likely to become excited with the prospect of yet another set of definitions amongst the many that already exist, the capabilities that result from a well-designed foundational taxonomy are significant.

Likewise, in order for our profession to evolve significantly, it is imperative that we operate with a common, logical, and effective understanding of our fundamental problem space. This Risk Taxonomy Technical Standard seeks to fill the current void and set the stage for the security profession's maturation and growth.

Note: Any attempt to describe the natural world is destined to be incomplete and imprecise to some degree due to the simple fact that human understanding of the world is, and always will be, limited. Furthermore, the act of breaking down and categorizing a complex problem requires that black and white lines be drawn where, in reality, the world tends to be shades of gray. Nonetheless, this is exactly what human-critical analysis methods and science have done for millennia, resulting in a vastly improved ability to understand the world around us, evolve, and accomplish objectives previously believed to be unattainable.

This Technical Standard is a current effort at providing the foundational understanding that is necessary for similar evolution and accomplishment in managing information risk. Without this foundation, our profession will continue to rely too heavily on practitioner intuition which, although critically important, is often strongly affected by bias, myth, and commercial or personal agenda.

2.1 What Makes this the Standard of Choice?

Although definitions and taxonomies already exist within the information security landscape, none provide a clear and logical representation of the fundamental problem our profession is tasked with managing – the frequency and magnitude of loss. For example:

- Existing taxonomies tend to focus on a subcomponent of the problem. Two current examples of work limited to particular areas of concern are the Common Weakness Enumeration (CWE) and the Common Attack Pattern Enumeration and Categorization (CAPEC).[1] However, while these two efforts are noteworthy, valuable, and consistent, most efforts are not consistent. In the absence of a common foundation it becomes difficult or impossible to tie together or interlink sub-taxonomies, which limits their utility to only the most narrow applications.
- Taxonomies are inconsistent in their use of common terms (e.g., "risk" in one taxonomy may translate to "vulnerability" in another). This makes normalization of data difficult if not impossible, and leads to confusion and ineffective communication, which can further erode credibility.

[1] Information about CWE is available at http://cwe.mitre.org, and information about CAPEC is available at http://capec.mitre.org.

- Documents that claim to describe "taxonomies" in fact provide definitions without clear or, in some cases, any descriptions of the relationships between elements. Absent these relationships, it becomes impossible to perform meaningful calculations even when good data is available.

The risk taxonomy described within this Technical Standard provides several clear advantages over existing definitions and taxonomies, including:

- There is a clear focus on the problem that management cares about – the frequency and magnitude of loss.
- Risk factor definitions are conceptually consistent with other (non-security) risk concepts that organization management is already familiar with.
- It enables quantitative analysis of risk through the use of empirical data (where it exists) and/or subject matter expert estimates.
- It promotes consistent analyses between different analysts and analysis methods.
- It provides a framework for describing how risk conclusions were arrived at.
- It effectively codifies the understanding of risk that many highly experienced professionals intuitively operate from but haven't had a reference for.
- It provides a reference and foundation for the evolution of specific sub-taxonomies.
- The multiple layers of abstraction within the model enable analysts to choose how deep/comprehensive they want to be in their analyses. This feature allows analysts to model risk in a cost-effective manner.

2.2 Who Should Use It?

This Technical Standard should be used by anyone seeking to:

- Understand how risk works and/or the factors that drive risk
- Consistently perform high quality risk analyses
- Develop or apply security metrics
- Evaluate, debate, or discuss the basis for risk conclusions
- Develop or apply risk analysis and assessment methodologies

A few examples of how the taxonomy can provide value are:

- Security organizations sometimes find that management rejects their risk conclusions and recommendations, in part because it's difficult to articulate the intuition and experience that led to those conclusions. The ability to explain how conclusions were arrived at using a logical and rigorous method can have a very significant impact on credibility in the eyes of management.

- Organizations often find that the quality and consistency of analyses performed by their security analysts vary widely. The Risk Taxonomy Technical Standard can be used to improve this by bringing everyone onto the same page with regard to terminology, definitions, and approach. This is especially helpful when bringing on staff who are newer to the profession, as it shortens the time it takes to make them effective.

- Metrics development and application are also improved by using the taxonomy to identify which data points are needed in order to support analyses, as well as where to get that data and how to use it. For example, data regarding threat contact frequency, the type of actions taken, which controls worked or failed to work, types and magnitude of loss, etc., can be extracted from incidents of all kinds (e.g., virus events, user errors, breaches, etc.) and used to support analyses.

- Organizations often engage external consultants to provide an impartial view of the organization's risk posture. The taxonomy can be used very effectively to evaluate the consultants' risk conclusions and recommendations, ensuring that findings aren't inflated (or underrated). This ability to more consistently and effectively analyze risk is a critical factor in enabling more cost-effective risk management.

2.3 Related Dependencies

In order to make effective use of this Technical Standard, risk assessment and analysis methodologies must provide data and/or estimates for each of the factors within the taxonomy. For example, if an assessment methodology leaves out or ignores threat event frequency, then conclusions resulting from the methodology will not align with the taxonomy nor will they faithfully represent risk.

Note that where empirical data doesn't exist for one or more of the risk factors, it is acceptable to use subject matter expert estimates. For practical purposes, quantitative estimates should not be precise. Instead, estimates should be provided as ranges (e.g., "a threat event frequency of 1 to 10 times per year") or as distributions (e.g., "minimum 1 time per year, most likely 7 times per year, with a maximum of 10 times per year") with some form of confidence rating that represents the level of certainty surrounding the estimates.

If qualitative estimates are used as inputs (e.g., "high", "medium", "low"), the estimates should ideally be mapped to a predefined set of quantitative ranges (e.g., "Medium = 1 to 10"). This enables the relationships between factors within the taxonomy to be represented mathematically, which enables more effective risk calculation. It also provides a means for comparison between analyses performed by different analysts (normalization), as well as a means of explaining how conclusions were arrived at.

If pure qualitative values are used (i.e., values that don't reference a quantitative range or distribution), then the taxonomy may be used as a structural reference rather than a framework for calculation.

Note that the decision to use qualitative or quantitative values should be driven by the needs and desires of those who will receive or base their decisions on the analysis results. A secondary factor that may drive this choice is whether the analyst is comfortable using quantitative estimates.

3 Risk Management Model

3.1 Risk Assessment Approach

All risk assessment approaches should include:

- An effort to clearly identify and characterize the assets, threats, controls, and impact/loss elements at play within the risk scenario being assessed
- An understanding of the organizational context for the analysis; i.e., what is at stake from an organizational perspective, particularly with regard to the organization's leadership perspective
- Measurement and/or estimation of the various risk factors
- Calculation of risk
- Communication of the risk results to decision-makers in a form that is meaningful and useful

3.2 Why is a Tightly-Defined Taxonomy Critical?

As alluded to earlier, without a logical, tightly-defined taxonomy, risk assessment approaches will be significantly impaired by an inability to measure and/or estimate risk factor variables. This, in turn, means that management will not have the necessary information for making well-informed comparisons and choices, which will lead to inconsistent and often cost-ineffective risk management decisions.

4 Functional

4.1 What is Defined?

This Technical Standard defines and describes the problem space our profession is tasked with helping to manage; i.e., risk. Each factor that drives risk is identified and defined. Furthermore, the relationships between factors are described so that mathematical functions can be defined and used to perform quantitative calculations.

4.2 What is In/Out of Scope and Why?

This Technical Standard is limited to describing the factors that drive risk and their relationships to one another. Measurement scales and specific assessment methodologies are not included because there are a variety of possible approaches to those aspects of risk analysis, with some approaches being better suited than others to specific risk problems and analysis objectives.

4.3 How Should it be Used?

This Risk Taxonomy should be used as a foundational reference of the problem space our profession is tasked with helping to manage; i.e., risk. Based on this foundation, methods for analyzing, calculating, communicating about, and managing risk can be developed.

Note that analysts can choose to make their measurements and/or estimates at any level of abstraction within the taxonomy. For example, rather than measure Contact Frequency, the analyst could move up a layer of abstraction and instead measure Threat Event Frequency. This choice may be driven by the nature or volume of data that is available, or the time available to perform the analysis (i.e., analyses at deeper layers of abstraction take longer).

5 Technical

5.1 Risk Taxonomy Overview

The complete risk taxonomy is comprised of two main branches: Loss Event Occurrence and loss magnitude. Within those two branches are the factors that drive the occurrence and magnitude of losses. Figure 2 lays out the higher-level abstractions within the framework.

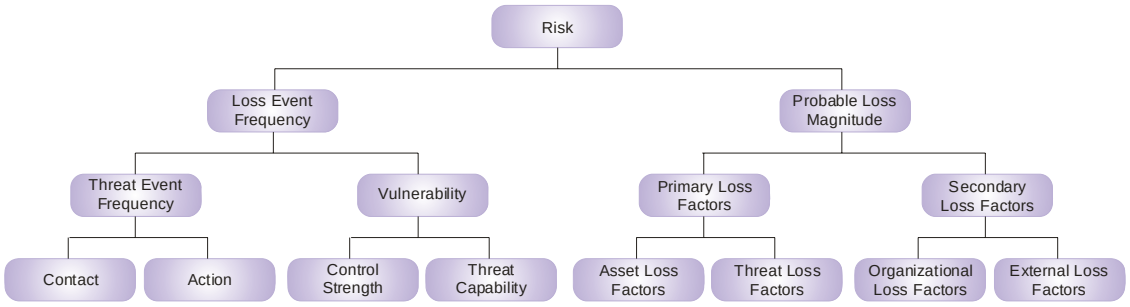

Figure 2

Note that this diagram is not comprehensive, as deeper layers of abstraction exist that are not shown. Some of these deeper layers are discussed further on in this document, but it is important to recognize that, theoretically, the layers of abstraction may continue indefinitely, much like the layers of abstraction that exist in our understanding of physical matter (e.g., molecules, atoms, particles, etc.). The deeper layers of abstraction can be useful in our understanding but generally aren't necessary in order to perform effective analyses.

Another point worth recognizing is that the factors within the Loss Event Frequency side of the taxonomy have relatively clean and clear cause-and-effect relationships with one another, which simplifies calculation. Factors within the Probable Loss Magnitude side of the taxonomy, however, have much more complicated relationships that defy simple calculation. As a result, loss magnitude measurements and estimates generally are aggregated by loss type (e.g., $xxx of productivity loss, plus $yyy of legal fines and judgments, etc.).

5.2 Component Definitions

5.2.1 Risk

Risk is the probable frequency and probable magnitude of future loss.

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

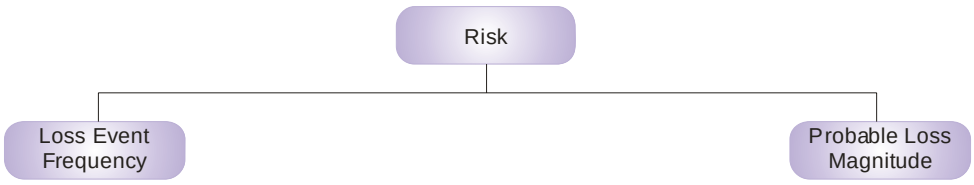

Figure 3

We will decompose the factors that drive Loss Event Frequency first, and then examine the factors that drive Probable Loss Magnitude.

5.2.2 Loss Event Frequency (LEF)

Loss Event Frequency (LEF) is the occurrence, 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 upon an asset, such that loss results. This leads us to our next two factors: Threat Event Frequency (TEF) and Vulnerability.

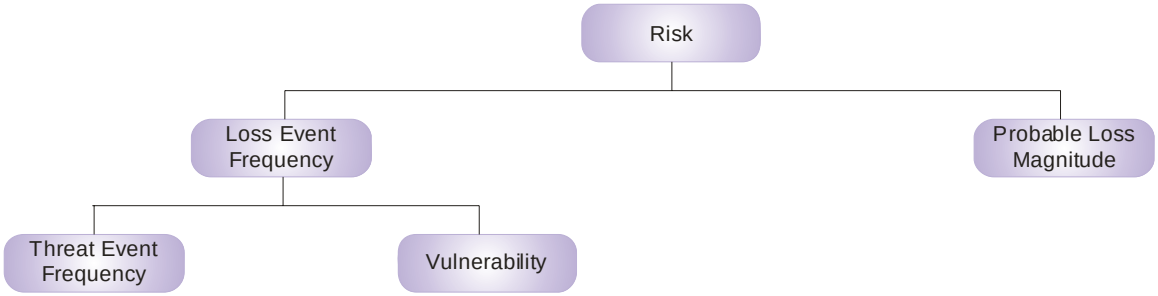

Figure 4

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

5.2.3 Threat Event Frequency (TEF)

Threat Event Frequency (TEF) is the occurrence, within a given timeframe, that a threat agent will act against an asset.

You will probably see the similarity between this definition and the definition for LEF above. The only difference is that the definition for TEF 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 5 adds these two factors to our taxonomy.

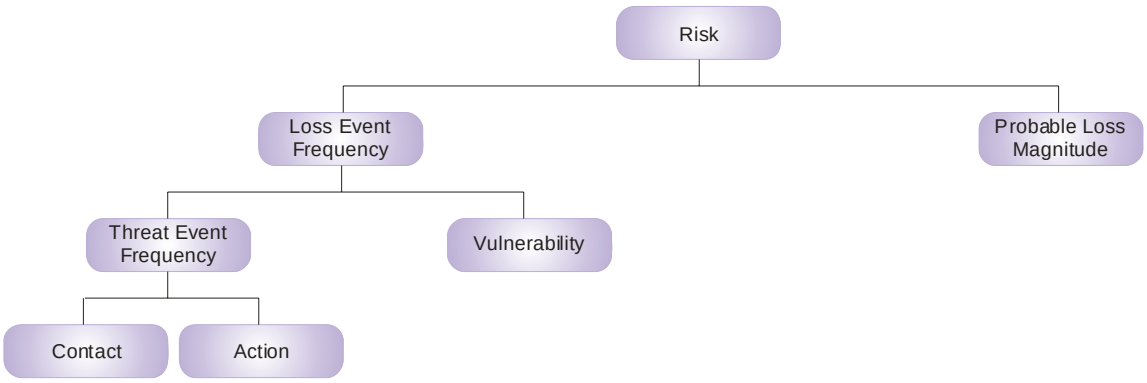

Figure 5

5.2.4 Contact

Contact is 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, as follows:

- **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. A useful analogy 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.

5.2.5 Action

Action is 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) and where the opportunity for decision on alternative loss causing and no loss exists.

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

- **Value** – the threat agent's perceived value proposition from performing the act.
- **Level of effort** – the threat agent's expectation of how much effort it will take to accomplish the act.
- **Risk** – the probability of negative consequences *to the threat agent*; for example, the probability of getting caught and suffering unacceptable consequences for acting maliciously.

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

Vulnerability exists when there is 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 (resistance capability). Figure 6 adds these factors to our taxonomy.

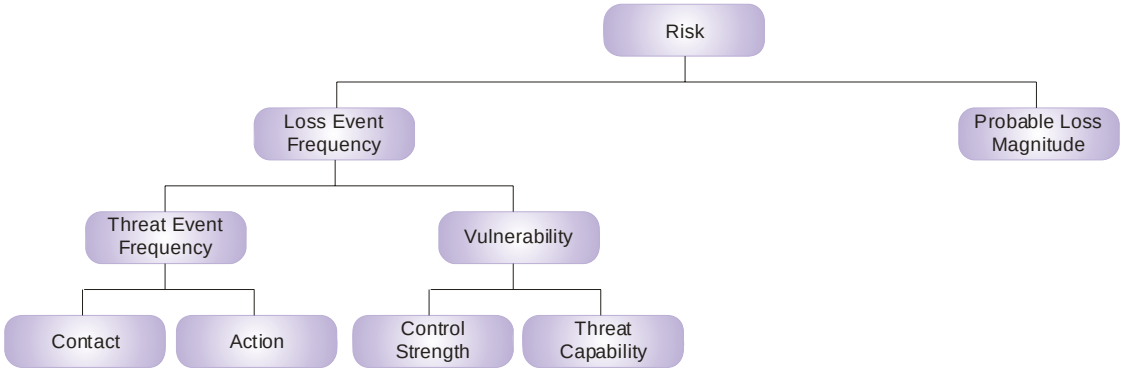

Figure 6

Vulnerability is always relative to the type of force and vector 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, a computer anti-virus 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 damage by 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.

5.2.7 Threat Capability

Threat Capability is the probable capability 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 is 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.

5.2.8 Control Strength (CS)

Control Strength (CS) is the strength of a control as compared to a baseline measure of force.

A rope's tensile strength rating provides an indication of how much force it is 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 the 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 is as well defined as PSI. Consider, however, password strength as a simple example of how we can approach this. We can estimate that a password eight characters long, comprised of a mixture of upper and lowercase 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.

5.2.9 Probable Loss Magnitude (PLM)

Probable Loss Magnitude (PLM) is the likely outcome of a threat event.

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, Probable Loss Magnitude (PLM) 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 uncertainty considerations altogether; we only cite the worst-case possibilities, or we try to be precise in our calculations. Excluding loss-related uncertainties from an analysis means that we are not analyzing risk (by definition, risk *always* has a loss component). Citing worst-case possibilities alone removes the probability 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. Management's 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 is difficult to evaluate loss probability; for example:

- It is 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.
- Frequent events are easier to treat probabilistically; rare or novel ones are hard.
- 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.

Information security incidents generally have a distribution that looks something like Figure 7.

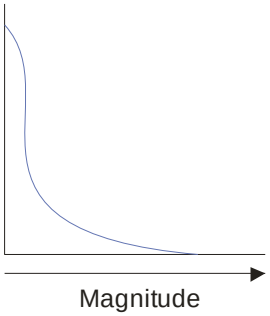

Figure 7

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.

5.2.10 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 can also 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 this Technical Standard, as follows:

- **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 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 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.

5.2.11 Loss Factors

All loss factors fall within one of the following four categories:

- Asset
- Threat
- Organization
- External

For reasons that will become clear, 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*.

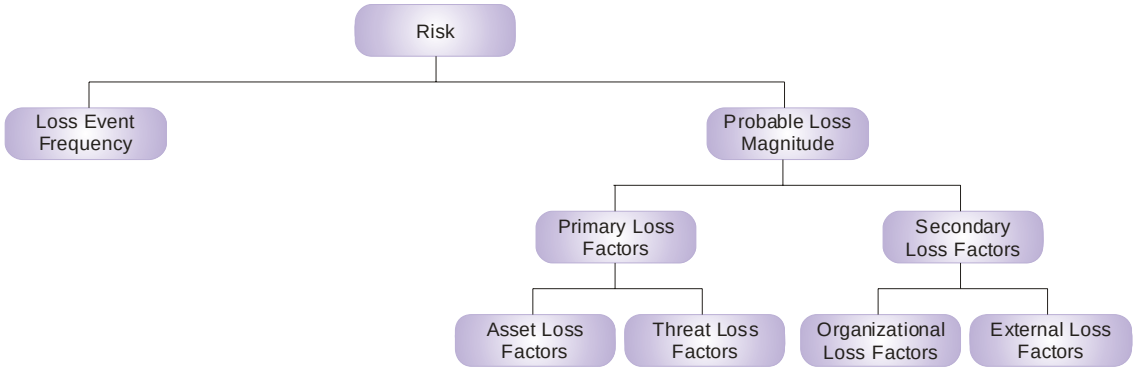

Figure 8

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 Technical Standard, we will limit our discussion to some of the most common and most important loss factors.

5.2.12 Primary Loss Factors

Asset Loss Factors

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

As 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** –the intrinsic value of the asset; i.e., the cost associated with replacing it if it has 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.

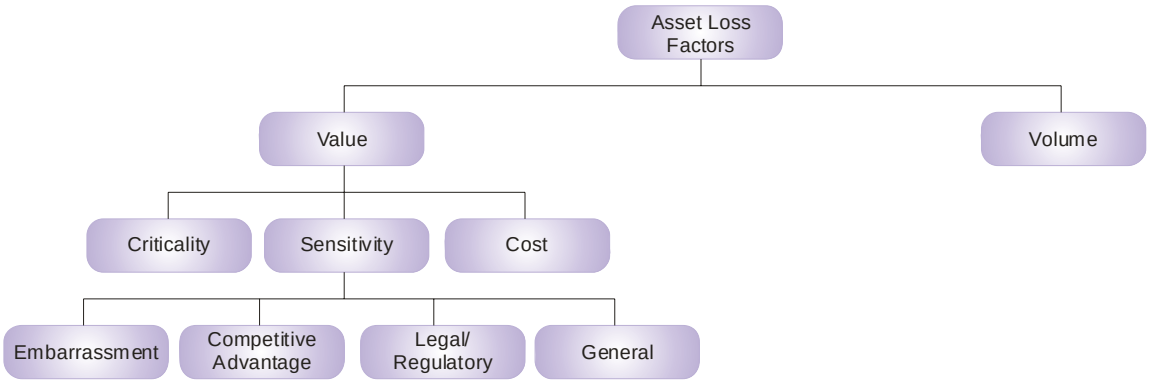

Figure 9

Asset volume simply recognizes that more assets at risk equals 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 threat 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 pornographic 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 is 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, kind of violation, and kind of exploitation of this violation 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. For this reason, it is critical to have a clear definition of your threat community in order to effectively evaluate loss magnitude.

This 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 capable of inflicting once the compromise occurs, while Threat Capability to Violate 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 this Technical Standard, 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 competencies*. We will not 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.

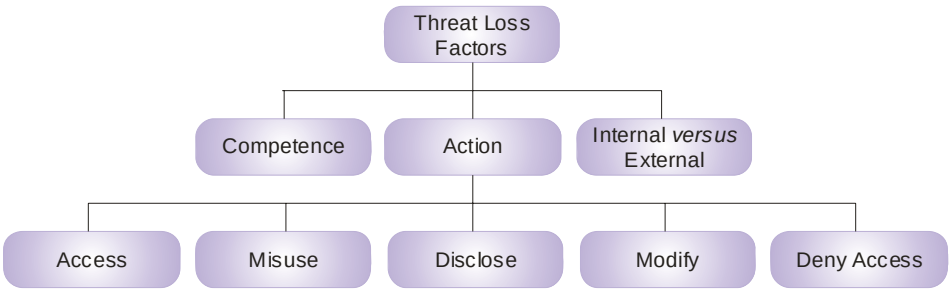

Figure 10

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 reputation losses because it is recognized that trusted insiders are exceedingly difficult to protect against.

5.2.13 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 will limit our discussion to four – timing, due diligence, response, and detection – as follows:

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 reputation 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 a response:

- **Containment** – 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** – an organization's ability to remove the threat agent; e.g., eradicating the worm.
- **Recovery** – 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 is critical to recognize, however, that response capabilities can also 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.

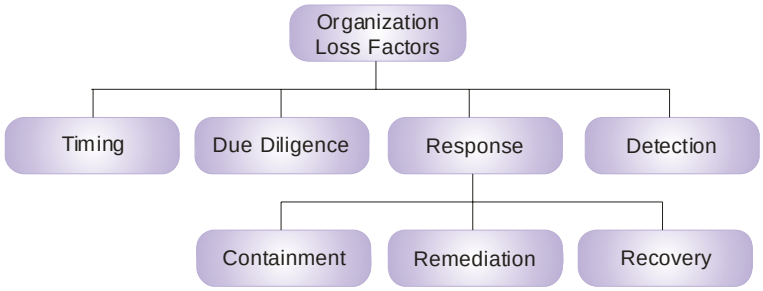

Figure 11

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 may still 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 five categories – detection, 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 on which all other external factors are predicated. 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 is required by law), or by accident.

The legal and regulatory landscape is primarily made up of three parts – regulations (local, state, federal, and international), contract law, and case law. Although this component of the external landscape is evolving rapidly, it is safe 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.

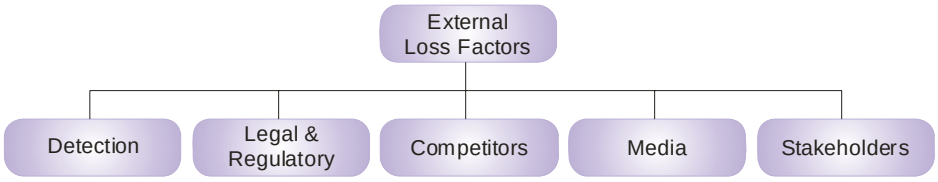

Figure 12

6 Example Application

This chapter provides an example of how the taxonomy may be used to perform a risk analysis. The analysis steps and charts shown are borrowed from the Introduction to Factor Analysis of Information Risk (FAIR). Note that any other well-designed analysis method could be used instead.

6.1 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 is entitled to use.

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

6.2 The Analysis: FAIR Basic Risk Assessment Methodology

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

- 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 Magnitude (PLM)
- Stage 4: Derive and articulate risk:
  - — Derive and articulate risk

6.2.1 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. A typical question in this scenario is 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 "they're all assets". In this case, however, we'll focus on the credentials, recognizing that their value is inherited from the assets they are 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
- Visitors to 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, we'll focus on the cleaning crew.

6.2.2 Stage 2: Evaluate Loss Event Frequency (LEF)

Estimate the Probable Threat Event Frequency (TEF)

Many people demand reams of hard data before they are comfortable with 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 are 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.

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 0.1 and 1 times per year
Very Low (VL)	< 0.1 times per year (less than once every 10 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)

Threat Capability (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 an SQL injection weakness – but scenarios like that don't lend themselves to an introductory document. In this case, all we're talking about here 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 an 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 Control Strength (CS)

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

Derive Vulnerability (Vuln)

Deriving vulnerability is easy once you've established your Tcap and CS. Recall from Section 5.2.6 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.

Tcap \ Vulnerability	VL	L	M	H	VH
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

Derive Loss Event Frequency (LEF)

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

TEF \ Vulnerability	VL	L	M	H	VH
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

In our scenario, given a TEF of Low and a Vulnerability of Very High, 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.

6.2.3 Stage 3: Evaluate Probable Loss Magnitude (PLM)

Using the previous seven steps, we have determined that the probability of a loss event in our scenario is Low (somewhere between 0.1 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 (organization 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, they 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, as follows:

- **Misuse** – employee records typically have information that can be used to execute identity theft, which introduces potential legal and reputation loss.
- **Disclosure** – employee records often have sensitive personal information related to medical or performance issues, which introduces legal and reputation 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 is 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.

	Loss Forms					
Threat Actions	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)	$10000000	–
High (H)	$1000000	$9999999
Significant (Sg)	$100000	$999999
Moderate (M)	$10000	$99999
Low (L)	$1000	$9999
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 is not worthwhile giving 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 on 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 reputation 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 reputation loss associated with customer flight and stock value could be Severe.

Note: Magnitudes will vary based on the size of the organization.

We are not 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 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 reduces the probability of a worst-case outcome.

In most cases it isn't worthwhile spending 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 is 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.

	Loss Forms					
Threat Actions	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)	$10000000	–
High (H)	$1000000	$9999999
Significant (Sg)	$100000	$999999
Moderate (M)	$10000	$99999
Low (L)	$1000	$9999
Very Low (VL)	$0	$999

Our rationale for these estimates includes:

- 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 reputation 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, as follows:

- 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.

6.2.4 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 0.1 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				
	Severe	H	H	C	C	C
	High	M	H	H	C	C
	Significant	M	M	H	H	C
PLM	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	Moderate
L	Low

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

A Risk Taxonomy Considerations

Extensive discussion in development of this Risk Taxonomy included considerations that can be grouped into four categories, as follows:

- Concerns regarding complexity of the model
- The availability of data to support statistical analyses
- The iterative nature of risk analyses
- Perspective

Many of these considerations are not so much critical of the FAIR framework, but rather are observations and concerns that apply no matter what method is used to analyze risk.

A.1 Complexity of the Model

There is no question that the proposed framework goes into greater detail than most (if any other) risk models. And, if usage of the framework required analyses at the deepest layers of granularity, then it would indeed be impractical for most risk analyses. Fortunately, most analyses can be performed using data and/or estimates at higher levels of abstraction within the model (e.g., measuring Threat Event Frequency rather than attempting to measure Contact Frequency and Probability of Action). This flexibility within the framework allows the user to choose the appropriate level of analysis depth based on their available time, data, as well as the complexity and significance of the scenario being analyzed.

Of course, the fact that the framework includes greater detail provides several key advantages:

- The aforementioned flexibility to go deep when necessary
- A greater understanding of contributing factors to risk
- The ability to better troubleshoot/critique analysis performed at higher layers of abstraction

Another consideration to keep in mind is that risk is inherently complicated. If it were not, then we would not need well-defined frameworks and we would not have challenges over analyzing it and communicating about it. Using over-simplified and informal models almost invariably results in unclear and inconsistent assumptions, leading to flawed conclusions, and therefore false recommendations. With that in mind, we recognize that even FAIR's detailed taxonomy isn't a perfect or comprehensive treatment of the problem. There are no perfect taxonomies/models of real-world complexity. It's just that we consider FAIR to be significantly more complete than what we're used to, and the best-analyzed and well-defined there is today.

With regard to communicating complex risk information to business decision-makers (who often want information like this delivered in simple form), the problem isn't inherently with the model but rather with the user. As is the case with any complex problem, we need to be able to articulate results in a way that is useful and digestible to decision-makers. It is also not unusual for management to ask how the results were arrived at. Experience has shown that having a rigorous framework to refer to in the explanation tends to improve credibility and acceptance of the results.

A.2 Availability of Data

In risk assessments, good data is especially difficult to acquire for infrequent events. In the absence of such data, how do we arrive at valid frequency estimates?

Good data has been and will continue to be a challenge within our problem space for some time to come. In part, this stems from the absence of a detailed framework that:

- Defines which metrics are needed
- Provides a model for plugging in the data so that meaningful results can be obtained

The FAIR framework has been proven in practice to help solve those two issues. It doesn't, of course, help us with those instances where data isn't available because events are rare. In those cases, regardless of what analysis method is chosen, the estimates aren't going to be as well substantiated by data. On the other hand, the absence of data due to the infrequency of events *is* data – of sorts – and can be used to help guide our estimates. As additional information is acquired over time, it is possible to adjust the initial estimates.

A.3 Iterative Risk Analyses

Due to the inherent complexity of risk, risk analyses tend to be iterative in nature. In other words, it is absolutely true that initial risk analyses tend to be "sighting shots" that often become more precise as additional analyses are performed. Furthermore, there comes a point of diminishing returns beyond which additional precision is not warranted given the necessary time and expense of deeper/broader analyses.

It is worthy of note that this observation is true of any analysis method, including the FAIR model.

A.4 Perspective

An alternative view held by some is that "exposure" should be the focus of our attention rather than "risk". The argument put forward here is that they consider "risk" to be the inherent worst-case condition, and "exposure" represents the residual risk after controls were applied.

Setting aside the possibility that those who hold this view misinterpret the definition of risk within the FAIR model, both issues are related (sort of a "before" and "after" perspective) and relevant. Fortunately, the FAIR framework provides the means to analyze both conditions by allowing the analyst to derive unmitigated risk as well as mitigated risk levels.

Glossary

Action

An act taken against an asset by a threat agent. Requires first that contact occurs 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.

Multi-level 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.

Index