SECURITY SERIES

Risk Management

The Open Group Guide

Van Haren Publishing

THE Open GROUP

Risk Management
The Open Group Guide

Other publications by Van Haren Publishing

Van Haren Publishing (VHP) specializes in titles on Best Practices, methods and standards within four domains:

- IT management
- Architecture (Enterprise and IT)
- Business management and
- Project management

Van Haren Publishing offers a wide collection of whitepapers, templates, free e-books, trainer material etc. in the **VHP Freezone**: freezone.vanharen.net

VHP is also publisher on behalf of leading organizations and companies:
ASLBiSL Foundation, CA, Centre Henri Tudor, Gaming Works, Getronics, IACCM, IAOP, IPMA-NL, ITSqc, NAF, Ngi, PMI-NL, PON, Quint, The Open Group, The Sox Institute

Topics are (per domain):

IT (Service) Management / IT Governance
ABC of ICT
ASL
BiSL
CATS
CMMI
CoBIT
ISO 17799
ISO 27001
ISO 27002
ISO/IEC 20000
ISPL
IT Service CMM
ITIL® V3
ITSM
MOF
MSF
SABSA

Architecture (Enterprise and IT)
Archimate®
GEA®
SOA
TOGAF®

Business Management
CMMI
Contract Management
EFQM
eSCM
ISA-95
ISO 9000
ISO 9001:2000
OPBOK
Outsourcing
SAP
SixSigma
SOX
SqEME®

Project/Programme/ Risk Management
A4-Projectmanagement
ICB / NCB
MINCE®
M_o_R®
MSP™
P3O
PMBOK® Guide
PRINCE2®

For the latest information on VHP publications, visit our website: www.vanharen.net, or freezone.vanharen.net for free whitepapers, templates and e-books.

Risk Management

The Open Group Guide

Colofon

Title:	Risk Management - The Open Group Guide
A Publication of:	The Open Group
Authors:	The Open Group
Editors:	Ian Dobson and Jim Hietala
Publisher:	Van Haren Publishing, Zaltbommel, www.vanharen.net
ISBN:	978 90 8753 663 3
Edition:	First edition, first impression, April 2011
Design and Layout:	CO2 Premedia bv, Amersfoort – NL
Copyright:	© The Open Group, 2011

For any further enquiries about Van Haren Publishing, please send an e-mail to: info@vanharen.net

Comments relating to the material contained in this document may be submitted to:
The Open Group
Apex Plaza, Forbury Road
Reading
Berkshire RG1 1AX
United Kingdom
or by electronic mail to:

ogspecs@opengroup.org

Preface

This book has been developed by **The Open Group,** 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.

Trademarks

Acknowledgements

Part 1: The Open Group Technical Standard: Risk Taxonomy

The Open Group gratefully acknowledges the contribution of:

- Alex Hutton, CEO, Risk Management Insight
- Jack Jones, CTO, Risk Management Insight

for contributing their FAIR (Factor Analysis of Information Risk) development work into the Security Forum of The Open Group, and their continued support in guiding the Security Forum members through The Open Group development and approval process to publish this Risk Taxonomy standard. The Open Group also acknowledges the members of its Security Forum who contributed to its development.

Part 2: The Open Group - Technical Guide
Requirements for Risk Assessment Methodologies

The Open Group gratefully acknowledges the contribution of:

- Alex Hutton, CEO, Risk Management Insight (www.riskmanagementinsight.com)

- Jack Jones, CTO, Risk Management Insight and the members of The Open Group Security Forum who contributed to its development.

Part 3: The Open Group Technical Guide FAIR–ISO/IEC 27005 Cookbook

The Open Group gratefully acknowledges the contribution of lead authors:
- Christopher Carlson, The Boeing Company
- Alex Hutton, Risk Management Insight, with the valued support of contributing author
- Anastasia Gilliam, Independent Consultant and the members of The Open Group Security Forum who contributed to its development.

References

The following documents are referenced in Part 1: The Open Group Technical Standard: **Risk Taxonomy**:

- 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.

The following documents are referenced in Part 2: The Open Group Technical Guide: **Requirements for Risk Assessment Methodologies**:

- COBIT (Control Objectives for Information and related Technology), Information Systems Audit and Control Association (ISACA); refer to www.isaca.org
- COSO (Committee of Sponsoring Organizations) Enterprise Risk Management Framework; refer to www.coso.org
- ISO/IEC 27002:2005: Information Technology – Security Techniques – Code of Practice for Information Security Management

- ITIL (Information Technology Infrastructure Library); refer to www.itil-officialsite.com/home
- OCTAVE (Operationally Critical Threat, Asset, and Vulnerability Evaluation); refer to www.cert.org/octave
- Risk Taxonomy Technical Standard, January 2009 (ISBN: 1-931624-77-1, C081), published by The Open Group
- FAIR - ISO/IEC 27005 Cookbook Technical Guide, November 2010 (ISBN: 1-931624-87-9, C103), published by The Open Group

The following documents are referenced in Part 3: The Open Group **Technical Guide FAIR–ISO/IEC 27005 Cookbook:**

- ISO/IEC 27005:2008: Information Technology – Security Techniques – Information Security Risk Management.
- ISO/IEC 27001:2005: Information Technoloy – Security Techniques – Information Security Management System – Requirements (ISMS)
- ISO/IEC 27002:2005: Information Technology – Security Techniques – Code of Practice for Information Security Management (Controls)
- Technical Standard: Risk Taxonomy (C081, ISBN: 1-931624-77-1), January 2009, published by The Open Group
- Technical Guide: Requirements for Risk Assessment Methodologies (G081, ISBN: 1-931624-78-X), January 2009, published by The Open Group

Contents

Introduction

This book brings together a set of three publications addressing risk management, which have been developed and approved by The Open Group. It is presented in three parts:

- Part 1: The Open Group Technical Standard for Risk Taxonomy
- Part 2: The Open Group Technical Guide to the Requirements for Risk Assessment Methodologies
- Part 3: The Open Group Technical Guide: FAIR – ISO/IEC 27005 Cookbook

Part 1: The Open Group Technical Standard for Risk Taxonomy

This part 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 part 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.

Part 2: The Open Group Technical Guide to the Requirements for Risk Assessment Methodologies

This part 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 intended audience for this part is anyone who is tasked with selecting, performing, evaluating, or developing a risk assessment methodology. This includes all stakeholders who have responsibilities covering these areas, including business managers, information security/risk management professionals, auditors, and regulators both acting as policy-makers and as law-makers.

Part 3: The Open Group Technical Guide: FAIR – ISO/IEC 27005 Cookbook
This part describes in detail how to apply the FAIR (Factor Analysis for Information Risk) methodology to any selected risk management framework. It uses ISO/IEC 27005 as the example risk assessment framework. FAIR is complementary to all other risk assessment models/frameworks, including COSO, ITIL, ISO/IEC 27002, COBIT, OCTAVE, etc. It provides an engine that can be used in other risk models to improve the quality of the risk assessment results. The Cookbook enables risk technology practitioners to follow by example how to apply FAIR to other risk assessment models/frameworks of their choice.

The primary target audience for this Cookbook is risk management analysts and practitioners, to help them to use ISO/IEC 27005 to achieve higher quality risk assessment results, especially given the lack of formal specificity in probabilism provided by ISO/IEC 27005, including its difficult appendices on creation of a probabilistic model.

PART 1 THE OPEN GROUP TECHNICAL STANDARD

Risk Taxonomy

Chapter 1 Introduction to risk taxonomy

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 the 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 whether the impact will 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 (Factor Analysis for Information Risk), 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. FAIR can be used 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 1.6) 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.

Chapter 2 Business case for a risk taxonomy

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 2.1 shows these chained dependencies.

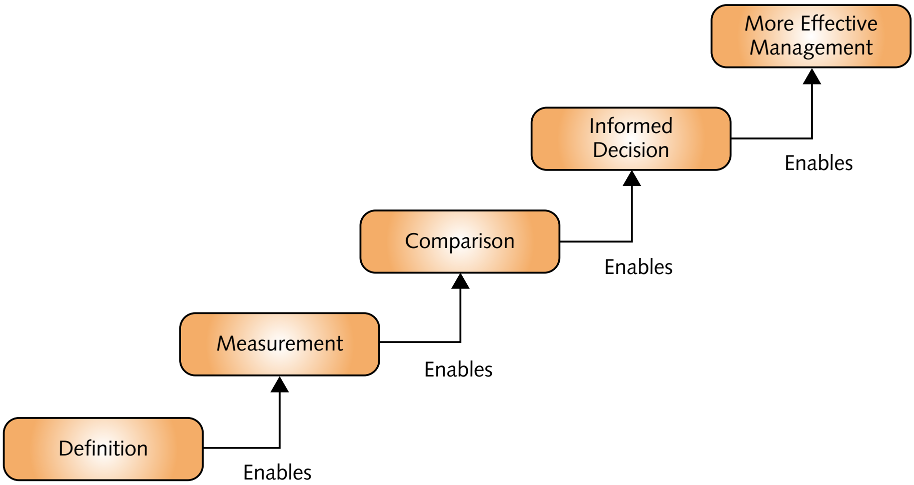

Figure 2.1: Risk dependencies

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 are 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.
- Documents that claim to describe "taxonomies" in fact provide definitions without clear descriptions (or, in some cases, without any descriptions) of the relationships between elements. Where information about these relationships is absent, 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.

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

- 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 this Technical Standard?

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 attitude to risk. 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.

Chapter 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.

Chapter 4 Functional aspects

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).

Chapter 5 Technical aspects

5.1 Risk taxonomy overview

The complete risk taxonomy is comprised of two main branches: Loss Event Frequency and Probable Loss Magnitude. Within those two branches are the factors that drive the occurrence and magnitude of losses. Figure 5.1 lays out the higher-level abstractions within the framework.

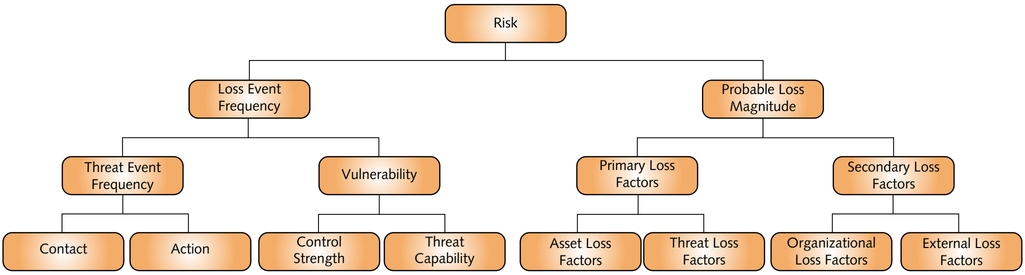

Figure 5.1: Risk taxonomy overview

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 Technical Standard, 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.).