

Levi Gundert

THE RISK BUSINESS

Second Edition

What Leaders Need to
Know About Intelligence
and Risk-Based Security

THE RISK BUSINESS

Second Edition

What Leaders Need to Know About
Intelligence and Risk-Based Security

By Levi Gundert

Foreword by Stu Solomon

With editorial assistance from Zane Pokorny



CYBEREDGE
P R E S S

The Risk Business: What Leaders Need to Know About Intelligence and Risk-Based Security, Second Edition

CyberEdge Group, LLC

1997 Annapolis Exchange Parkway

Suite 300

Annapolis, MD 21401

(800) 327-8711

www.cyber-edge.com

Copyright © 2023, CyberEdge Group, LLC. All rights reserved. Definitive Guide™ and the CyberEdge Press logo are trademarks of CyberEdge Group, LLC in the United States and other countries. All other trademarks and registered trademarks are the property of their respective owners.

Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, without the prior written permission of the publisher. Requests to the publisher for permission should be addressed to Permissions Department, CyberEdge Group, 1997 Annapolis Exchange Parkway, Suite 300, Annapolis, MD, 21401 or transmitted via email to info@cyber-edge.com.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on CyberEdge Group research and marketing consulting services, or to create a custom *Definitive Guide* book for your organization, contact our sales department at 800-327-8711 or info@cyber-edge.com.

ISBN: 978-1-948939-35-5 (paperback)

ISBN: 978-1-948939-38-6 (eBook)

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgements

CyberEdge Group thanks the following individuals for their respective contributions:

Copy Editor: Susan Shuttleworth

Graphic Designer: Debbi Stocco

Production Coordinator: Jon Friedman

Foreword



It is my pleasure and honor to recommend this work as essential reading for both technically astute and business-minded security practitioners. I believe it will help us as we work as a community to protect the availability and integrity of our systems, data, and operations in the face of malicious cyber actors.

For nearly 20 years, I have watched our security industry struggle to strike the right balance between the technical tools of our trade — usually complex, sometimes elegant, and often expensive — and the practical outcomes we try to achieve. As is so often the case in our industry, the predominant topic of discussion is technical controls and tools. This focus, while necessary (and let's face it, the source of great job creation for ourselves and so many of our colleagues and friends), actually prevents us from embracing a holistic approach to intelligence.

This approach requires an evolution toward risk-based cybersecurity, which is predicated on the ability to express the value of security activities in terms of measurable and defined outcomes based on risk reduction. A holistic approach to intelligence also requires a rich understanding of the threat environment, a clear appreciation of the concept of criticality, and an awareness of the potential impact of cyberattacks on business operations.

In this updated second edition, my close friend and colleague Levi Gundert deftly bridges the chasm between technology and focused risk reduction. He describes how to create an environment where operational risk is identified and managed down to an acceptable level. Levi reviews core concepts of traditional intelligence and describes advanced techniques that can be

used to identify and quantify threats based on adversary activity, intent, and capabilities. The end result is a clearer picture of the risks posed to enterprises by threat actors.

I encourage you to use this book as a practical guide and apply its ideas to your strategic and operational security challenges. Its lessons will help you break down problems into bite-sized chunks, enrich your understanding of the threat environment, make decisions with business criticality and operational context in mind, and take actions that are measured and focused on risk reduction.

Stu Solomon

President — Recorded Future
Charlotte, NC

Preface to Second Edition

The first version of *The Risk Business* was produced mainly to inform readers about the benefits of quantifying cyber risk and the critical role of intelligence in that quantification process. Threat Category Risk (TCR) was presented as a practical and easy to implement alternative to the FAIR (Factor Analysis of Information Risk) model.

After three years and hundreds of conversations with chief information security officers (CISOs) in literally every industry vertical and geography, I came to realize that although risk quantification is achievable, gaining acceptance of results at the board level is challenging. Enterprise risk management (ERM) groups have been describing and measuring many types of business risk for decades, but cyber risk is relatively new. Executives still prefer to communicate in general terms that center around “likelihood” and “potential impact.”

But, although winning hearts and minds for cyber risk quantification is a long-term effort, CISOs (and other leaders) need to start the conversation now. And you can do that in ways that are rigorous and build on the foundation of cyber risk quantification, and also compel the attention of executives and board members.

In this second edition of *The Risk Business*, I keep the foundation of TCR and risk quantification, but focus on helping security leaders tell their story through the Intelligence to Risk (I2R) Pyramid, the five risk impacts taxonomy, second-order thinking, and other techniques. I also describe how good

intelligence can guide and enhance risk storytelling. Whether you are attempting to quantify cyber risk, or simply to communicate the business value of your security controls and processes, I believe this book will provide you with new ideas and effective tools.

Levi Gundert

Contents

Foreword	iii
Preface to Second Edition	v
Acknowledgements	x

Section 1: Risk-Based Security and Risk Impacts

Chapter 1: The Case for Risk-Based Security	3
What This Book Covers	4
Why You Should Listen to Me.....	5
Risk Is the Language of Business	6
Threat-Driven Security Programs.....	7
Compliance-Driven Security Programs	10
FUD and Herd Mentality	14
Risk-Based Security and Communication with Top Management ..	15
Chapter 2: Risk Impacts	19
Risk Impacts and Adversary TTPs	21
Reducing Uncertainty with Intelligence	22
Chapter 3: Legal or Compliance Failure	27
More Regulations, More Failures, Rising Costs of Non-Compliance	28
Data Privacy: The New Legal Frontier.....	30
Legal Failures	31
How Intelligence Can Prevent Legal and Compliance Failures	32
AI Advances and Cybercrime Advantages	33
Chapter 4: Operational Disruption	37
Disrupting Networks and Services.....	39
Destroying Data and Disabling Systems.....	39
Threats to Operations Are Evolving Rapidly	40
Chapter 5 : Brand Impairment	43
Ideology and Cyber Targeting	48
Chapter 6: Financial Fraud	51
Chapter 7: Competitive Disadvantage	57

Section 2: How to Quantify Risk

Chapter 8: Risk — to Quantify or Qualify? 63

 The Meaning of Risk.....63

 Defining Risk as Monetary Loss.....64

 But We *Can’t* Estimate Risk in Cybersecurity (Can We?)66

 The Power of Estimation.....66

 Estimation Training Is Important67

 Monte Carlo Simulations 68

 After Quantification, Communicate with Stories69

 A Last Resort 71

Chapter 9: Quantifying Risk with the Threat Category Risk Framework..... 73

 Introducing the Threat Category Risk Framework74

 The TCR Threat Categories75

 Walking Through an Example: Credential Reuse 77

 Running the Monte Carlo Simulation.....79

 Estimating the Value of New Security Controls82

 Complementary to Control Validation.....82

Chapter 10: Updating the Framework: RTDs and Threat Intelligence..... 83

 Relevant Threat Deltas (RTDs).....83

 Why Threat Intelligence Is Critical for Risk-Based Cybersecurity.. 84

Section 3: Intelligence and Risk Management for Business

Chapter 11: Strategic Threat Intelligence..... 89

 The Value of Intelligence..... 89

 What Is Intelligence?..... 90

 Intelligence Leads to a Persistent Decision Advantage 90

 Start by Strengthening Basic Security Controls..... 91

 Don’t Rely on Daily Threat Reports..... 91

 Daily Threat Reports Versus Useful Reports92

 Don’t Create Reports for Nonexistent Audiences.....94

 Stop Saying “Actionable”.....95

 Risk-Based Analysis Helps Threat Intelligence Produce Operational Outcomes95

 Sourcing.....97

 Private Sector Sources.....99

 Staffing and Community Support100

Avoid Siloing.....	102
Tooling and Measurements.....	103
Chapter 12: The Intelligence to Risk (I2R) Pyramid and Additional Considerations for Strategic Threat Intelligence	105
I2R Pyramid Example: Ransomware.....	109
Three As for Addressing New TTP Instances	111
Security Control Validation.....	112
Workflows and Outcomes	113
Attribution	118
Chapter 13: Operational Threat Intelligence and the Outcomes Matrix.....	119
What Is Operational Threat Intelligence, and Why Do I Need It?..	119
Good and Bad Indicators	121
Measurement.....	122
The Outcomes Matrix.....	124
Chapter 14: Five Critical Cybersecurity Functions.....	127
Security Starts with Preventing, Detecting, and Removing Unauthorized Access	127
The Problem of Identity and Access Management.....	128
IAM That IAM	131
The Role of Intelligence in Vulnerability Management.....	132
Third- and Fourth-Party Risk	133
Email Security	135
Cloud Security	137
Successful Business Is Risk Management	138

Acknowledgements



Version 0.1 of this book was largely written with T-Rex arm posture on a tray table somewhere north of 30,000 feet while crossing the United States or a large stretch of ocean. This second edition was primarily hammered out on hotel room desks. I'm grateful to my whole crew at home for their patience with me and for inspiring me to hustle and grind every day. None of this would be possible without them.

Throughout my career I've worked for incredible organizations stocked with smart and talented people. Recorded Future stands out for its exceptional, continual pursuit of solutions to customer problems. I'm enormously blessed to work daily with Futurists (my colleagues at Recorded Future) while attempting to contribute to the world's best intelligence platform. Thank you, Recorded Future, for all the memories and for your assistance with this effort.

I'm especially grateful to Dylan Davis, Zane Pokorny, and Lucas Clauser for their energy, input, and editorial attention while turning this second effort into a real book.

The images at the beginning of each chapter were created with AI engines available at prompthunt.com.


Finally, this book would not be possible without the considerable influence and review of industry titans and all my amazing Recorded Future colleagues, who helped shape my thinking.

Levi Gundert

Section 1: Risk-Based Security and Risk Impacts

Chapter 1

The Case for Risk-Based Security



Building a successful cybersecurity program isn't easy. One of the key factors is how you define success. In fact, if you define success the wrong way, you will end up with:

- Poor allocation of resources, including time
- Misleading metrics that create the wrong incentives
- Grave failures of communication between security leaders and business executives

That is exactly the situation in which many cybersecurity organizations find themselves today. Their cybersecurity programs are either threat driven and focused on deploying industry best practice security controls to meet the latest cyber threats, or compliance driven and organized to “check the boxes” on security and privacy requirements produced by third-party standards organizations. Both approaches have flaws.

I define success in cybersecurity as a material and measurable reduction in cyber risk (which is considered part of operational risk) *and* a persistent decision advantage over adversaries. Adopting my definition will drive you to assemble processes and tools that lead to better allocation of resources, meaningful metrics that drive the right incentives, and productive discussions between IT security professionals, executives, and line managers.

This book is designed to help you achieve those objectives.

What This Book Covers

In the first section (Chapters 1-7), I will make the case for risk-based security and contrast it with threat-driven and compliance-driven cybersecurity programs. In addition, we will review the five leading types of risk impact: legal or compliance failure, operational disruption, brand impairment, financial fraud, and competitive disadvantage.

In the second section (Chapters 8-10), I will describe a process for quantifying risks in monetary terms, so you can prioritize new security controls and communicate the value of risk reductions in language non-technical managers can understand. This process is far simpler, faster, and more accurate than you would expect.

In the third and final section (Chapters 11-14) we will focus on intelligence and how it strengthens risk reduction in particular and cybersecurity programs in general. I will also share insights into managing an intelligence program and integrating it with five critical cybersecurity functions.

If you are an information security practitioner struggling to relate to your business, this book is for you. If you're an executive looking to make savvy security decisions based on strong risk metrics, this book is for you. This book will help you create a persistent advantage for better security so your business can focus on being profitable.

Why You Should Listen to Me

Before we dive into these topics, I hope you'll allow me to indulge in a little reminiscing as I describe my background in the field. I don't want you to think I'm a joker off the street calling on you to upend your entire security strategy.

My first thoughts around the role that risk reduction plays in business strategies came when I was in university. I remember reading a book by Eliyahu M. Goldratt called "The Goal" in an operations management class. The book's message — which is somewhat counterintuitive in our age of companies hyper-focused on revenue growth — is that profitability is the only meaningful business goal. For a business to thrive in perpetuity, every employee should be focused on this one bottom-line goal of increasing profits.

In the early 2000s, my work as a network security administrator gave me a front-row seat to many cyber events disrupting operations at healthcare and financial services companies. Some analysts hypothesized that IT system interruptions were contributing to decreased productivity, resulting in lost revenue, but no one ever quantified the loss.

Fast-forward a few years. I was sporting a badge and gun while pursuing cybercriminals around the world as a member of the U.S. Secret Service's electronic crimes task force. I quickly realized that the concept of threat intelligence was critical to criminal investigations, aiding in suspect attribution and successful prosecutions. My successful cases started with proactive intelligence collection, almost always in coordination with brilliant minds in the private sector.

It wasn't long before I rejoined the private sector (no more flying armed, but better data). Between consulting for clients and contributing to the defense of an enterprise, I realized that a specific articulation of risk was the greatest challenge facing senior security and business leaders. Additionally, I recognized that although intelligence is now a critical capability for risk management, the private sector continues to struggle with defining the value of outcomes for intelligence teams and their workflows.

Risk Is the Language of Business

Cybersecurity professionals tend to see themselves as business enablers. As defenders, they keep the bad guys out so that the business can operate uninterrupted.

However, the C-suite and board of directors are more concerned with profitability. Often, those at the top of the organization see cybersecurity groups as cost centers dragging down the bottom line. Changing that cost center perception is critical to building a successful cybersecurity program.

Someone once said, “There should only be two types of people in a business — those who make things, and those who sell things.” Today, there is a third category: those who defend things. This category is as necessary as the other two. However, while we have widely accepted procedures and metrics for measuring how making things and selling things contribute to the profitability of the enterprise (indeed, we have large accounting organizations set up to do exactly that), most organizations have barely started to think about how to measure the contribution of people who defend things.

How do you measure and communicate the value of a basic security control action? The answer lies in the language of risk. Senior decision makers don’t necessarily understand the language of security or even technology, but they speak the language of risk.

As a cybersecurity professional, your goal should be to quantify, or at least to clearly qualify, how every potential cybersecurity investment in staff and tools will reduce risk.

If you can do that, you will find it much, much easier to:

- Set priorities among possible cybersecurity investments based on real outcomes for the enterprise
- Justify budget requests for each investment and for the overall level of investment in cybersecurity
- Work productively with executives and line management to estimate risk and find the most cost-effective ways to reduce it

Hard work and smart choices are required to achieve these goals, but it can be done. We will discuss many techniques throughout this book, but first, let's explore the problems that occur when you build your cybersecurity program around threats or compliance requirements.

By the way, in the context of risk analysis, to “qualify” means to classify by level. Typically, this means dividing risks into low, medium, and high categories based on criteria such as probability of occurrence and magnitude of impact. I fundamentally disagree with an approach that uses subjective labels as its primary metric because such labels are open to multiple interpretations. This is why “quantifying” risk is superior to “qualifying” risk.

Threat-Driven Security Programs

Threat-driven security programs implement industry best practice security controls based on the latest evolution of cyber threats. Little thought is given to whether a new category of threat poses a risk to the defender's business.

The distinction between threats and risks is extremely important. Threats are dangers — an unrealized possibility that could potentially harm your organization. But not every cyber threat is a risk. If an existing control (process, technical, or otherwise) can defeat the threat, then it is not a risk for you. If controls are insufficient or absent, then a threat can quickly become a risk.

For example, Hancitor (also known as Chanitor and TorDal) is a label for malicious code (malware) that acts as a trojan capable of downloading additional trojans. Hancitor is a payload typically delivered by email. When it first surfaced in 2014, what made it worthy of attention was its ability to perform process hollowing: injecting code into a legitimate running process to disguise it from endpoint security software (like an antivirus client).

Chanitor (Hancitor, Tordal) – Malware
Actions

- 3 Analyst Notes
- 5 Insikt Group Notes
- 10 000+ References to This Entity
- First Reference Collected on Oct 30, 2014
- Latest Reference Collected on Apr 15, 2019
- ★ Curated Entity
- Malware Category Trojan

Show recent cyber events involving Chanitor in [Table](#)

Show all events involving Chanitor in [Table](#)

Threat Research from Insikt Group

All 5 | [Indicator 1](#) | [Cyber Threat Analysis 1](#) | [Flash Report 3](#)

All 5 | [Primary 4](#) | [Related 1](#)

eFax-Themed Hancitor Malspam Campaign Indicator

The indicators attached to this Insikt Note are reportedly linked to an eFax-themed Hancitor (aka Chanitor or Tordal) malspam campaign from email address [efax@redelephantpizza.com](#). [Full note](#)

Source Insikt Group on Aug 21, 2018, 04:00 • [Note Actions](#)

New Phishing Campaign Uses Lure of Free Air France Tickets [Cyber Threat Analysis](#)

Delta Receipt-Themed Spam Campaign Delivers Hancitor Malware [Flash Report](#)

HSBC Loan-Themed Spam Delivers Hancitor Malware [Flash Report](#)

Recorded Future Intelligence Card for Hancitor/Chanitor (snapshot taken on April 15, 2019)

The author(s) of Hancitor innovated when they designed the code to install itself on victim machines as surreptitiously as possible.


If you weren't sure whether your endpoint security controls, such as your antivirus software or endpoint detection and response (EDR) client, were capable of detecting Hancitor's process hollowing, or you determined that there was a gap in your controls, then Hancitor posed a risk to your business.

However, if your endpoint security controls were capable of detecting process hollowing, then Hancitor was a threat *but not a risk* for your organization.

A more recent example of threat that might or might not be a risk is "Roasting oktapus," which was documented by Group-IB. Threat actors compromised Okta multi-factor authentication (MFA) credentials by social engineering attacks on employees at multiple companies via SMS (text messaging).

The cleverly planned attack was expertly executed in less than an hour. When the digital dust settled, thousands of employee

credentials across multiple companies were used to gain unauthorized access. The attack was successful because the Okta MFA controls that had been implemented were vulnerable to confused employees who responded to well-crafted SMS messages.



Flow of the Roasting Oktapus coordinated attack (Source: Recorded Future)

According to reports in the press, one company, Cloudflare, was immune to the attack. Why? Cloudflare didn't leave email security up to its employees. The company implemented a stronger form of MFA that required a hardware token (Yubikey) to authenticate. It didn't matter that threat actors could target Cloudflare employees with social engineering attempts because those actors were never going to be in physical possession of employee tokens. In other words, Roasting oktapus was a threat for everyone, and a risk for most, but it was not a risk for organizations like Cloudflare that had put the right controls in place.

Cybersecurity professionals should never act on a threat before understanding whether it represents a risk to the business. The effort may waste valuable resources (time and money). The difference may seem academic, but the practical application of this philosophy is critical to a cybersecurity program's success.

Threat-driven security programs expend resources on threats that are not actual risks, or are only minor risks, and miss opportunities to address serious threats. For example, security organizations that measure and reward teams for the number of threats mitigated create an incentive to work on threats that can be fixed quickly, even if they pose little risk, and to neglect complex and potentially more costly threats. IT managers who request funds to protect against the threats that are in the headlines, rather than issues that pose imminent risks to the enterprise, are more likely to suffer major data breaches.

Compliance-Driven Security Programs

The goal of compliance-driven security programs is to increase the organization's maturity level as measured by criteria published by a third-party standards organization. For example, ISO 27002 and the National Institute of Standards and Technology (NIST) Cybersecurity Framework (CSF) are comprehensive compliance frameworks that periodically revise best practice guidance. Businesses following compliance-driven security programs rely on these third-party organizations to accurately identify threats and provide guid-

ance on remediation actions for each category of threat.

These compliance frameworks are helpful guidelines, but ambiguity around emerging technologies and infrequent updates can leave gaps in their requirements.

Third-party compliance frameworks do encourage risk measurements, but they don't provide prescriptive guidance about how to perform these measurements. The primary way to communicate value from a compliance-led cybersecurity program is to announce when a new maturity level is reached. However, not only is that metric subjective, it's an unreliable indicator of risk.

Target Corporation is an often-cited textbook example of a devastating data breach suffered by an organization that measured very well on compliance.¹ In 2013, Target was certified as PCI (Payment Card Industry) compliant. But the initial unauthorized access originated from a third-party heating, ventilation, and air conditioning (HVAC) vendor, and at that time PCI compliance didn't require continual third-party access and risk auditing.² Achieving compliance with the framework didn't prevent the breach.

Compliance-led security programs are dangerous for two reasons:

1. They produce a “check the box” mentality and encourage an attitude of fulfilling the letter, but not the spirit, of the law.
2. Today, businesses continuously introduce new technologies that increase complexity and risk, and the standards organizations can't keep pace.

Let's look at these two points in greater depth.

Following a compliance-based security program by just checking a series of boxes may lead to complacency once a few best practices have been implemented. Governance and compliance obligations must be fulfilled, but compliance frameworks should be used as a tool, not as the end goal or mission.

For example, let's say I tell you it's a best practice to build a fence around your home to keep intruders out. If you build a

1. <https://www.technewsworld.com/story/80160.html>

2. <https://www.bankinfosecurity.com/target-update-a-6489>

two-foot-high fence made of popsicle sticks you have fulfilled the letter of my request, but not the spirit. Are we just keeping rabbits out of the garden, or are we trying to stop thieves from breaking into the property?




If a compliance framework (like COBIT, ISA, HIPAA, or PCI) mandates that you deploy stateful inspection firewalls, and you comply but mistakenly configure the rules to allow all incoming traffic, then you have satisfied the requirement but haven't reduced risk for the organization. Even worse, you have wasted resources on an ineffective solution.

If that same compliance framework requires a 24/7 security operations center (SOC) to manage alerts and you outsource the job to an incompetent vendor, then once again you've checked the box, but you've made the organization worse off by both wasting resources and providing a false sense of security.

When businesses adopt new technologies, they create new opportunities but also increase risk from cyber threats. New technologies add complexity to already complicated environments. Adversaries love complexity because it increases opportunities for successful attacks. More systems, more vendors, more suppliers, and less control of data mean the traditional security architecture playbooks must be revised to reflect a world where threats from third- and fourth-party integrations pose greater risks. Standards organizations simply can't move fast enough to address these new challenges, so their prescriptions will never cover all of the new risks.

If you're not convinced that compliance frameworks fall short as an overarching goal, consider the increasingly dire news headlines around data breaches from 2016 to 2023. In 2016, Yahoo confirmed a data breach affecting 500 million customer accounts. In 2017, Equifax (one of three major American credit reporting agencies) sustained a data breach resulting in the theft of PII from roughly 150 million American citizens. In 2018, Starwood Hotels and Resorts announced that it had been victimized for years, leading to the theft of the PII for around 500 million guests. In 2021, ransomware knocked systems offline at Colonial Pipeline, which resulted in widespread fuel outages on the East Coast of the United States. The list is endless. A Recorded Future search for 2022 data breaches returned 400,000 results. The scale of the problem is difficult to grasp.



Timeline of data breaches in 2022, showing some of the more than 400,000 results (Source: Recorded Future)

I am not denying that compliance is a critical goal. It has become even more important and more challenging recently. Regulations are proliferating in almost every industry and geography. Organizations can no longer assume that matching the controls competitors have implemented will provide a safe harbor in the event of a breach. Instead, they need to perform a rigorous analysis of risk and compliance costs against the backdrop of new domestic regulations, such as the California Consumer Privacy Act (CCPA), as well as international mandates like the European Union's General Data Protection

Regulation (GDPR) and the European Union Agency for Cybersecurity (ENISA) NIS Directive 2. Regulatory agencies are beginning to increase due diligence expectations and impose significant financial penalties on businesses for control failures.

However, the fact that audit compliance is a key domain for success does not mean it should be the principal guide for balancing risk against security costs and investments.

FUD and Herd Mentality

I would like to call out two other concepts that can be useful at times but should never be allowed to override a systematic analysis of risk.

Fear, uncertainty, and doubt (FUD) can be a very effective motivator. There is nothing like a data breach at a competitor to galvanize executives and boards into authorizing funds for needed security controls. However, security leaders should use FUD sparingly when communicating with top management because:

- The effects of FUD are often temporary
- Threats to competitors may not represent the biggest risks to your organization
- Overuse of FUD creates mistrust of the security team among executives and boards of directors, undermining the confidence needed for good, risk-based decisions

Comparisons with similar organizations can also be useful in the right situations. Every executive wants to hear that their security program is performing on par with competitors'. Parity is easy to understand and can be translated into concrete requests. For example, a chief legal officer of a publicly traded company recently confided to me that the security budget process boils down to a competitive analysis. If a new control is being implemented by the competition, then the business justification is strong.

However, while achieving parity with competitors might help avoid a future regulatory penalty based on a "reasonable security" review, the herd mentality doesn't translate into

savvy risk management. Every company is different. Business and security priorities should align through risk management strategy that is uniformly communicated at every level of an organization.

Risk-Based Security and Communication with Top Management

In the executive summary, the World Economic Forum's Global Cybersecurity Outlook 2023 made a few notable observations that reinforce the recognized need for tighter communication between cyber and business leaders:

- Structured interactions between cyber and business leaders are becoming more frequent — 56% of security leaders now meet monthly or more often with their board. This is rapidly narrowing the cybersecurity perception gap. However, more needs to be done to promote understanding between business and security teams to support effective action by organizational leaders.
- Building a security-focused culture requires a common language based on metrics that translate cybersecurity information into measurements that matter to board members and the wider business.
- Ultimately, cyber leaders must present security issues in terms that C-level executives can understand and act on. Business leaders, for their part, need to accept more accountability for operational cyber requirements to advance their organizations' overall cyber capabilities.³

In 2022, my colleagues Anna Iskenderian, Jesse Nuese, and Jakob Wolk and I compiled data on more than 400 public cybersecurity failures that led to some type of financial loss.⁴ We found that the median loss amount per event across all companies analyzed was \$1.5 million.


3. https://www3.weforum.org/docs/WEF_Global_Security_Outlook_Report_2023.pdf

4. <https://go.recordedfuture.com/hubfs/reports/managing-cyber-risk-stakeholder-capitalism.pdf>

That sounds like a nice statistic everyone can use to understand the impact of a loss event: “The median financial loss for a cybersecurity failure is \$1.5 million.”

But we also found that financial loss varied widely across different types of events, ranging from just over \$1 million for PII exposure to \$3.22 million for the theft of trade secrets. Moreover, median loss varied enormously by industry. The median loss per event was \$15.2 million for industrial companies, 50 times the median loss of \$300,000 for communication services organizations.

Depending on your industry and the type of loss event you are discussing, the simple median figure of \$1.5 million can be extremely misleading.



Median financial loss per event, by loss type (Source: Recorded Future)

Industry	Median of Loss Amount (\$)	Count of Industry
Industrials	\$15,200,000	22
Energy	\$10,600,000	5
Consumer Staples	\$4,662,500	17
Information Technology	\$4,625,000	54
Finance	\$2,750,000	63
Government	\$2,700,000	3
Real Estate	\$1,500,000	2
Consumer Discretionary	\$1,339,481	41
Healthcare	\$1,020,000	117
Services	\$457,059	74
Communication Services	\$300,000	15

Median financial loss per event, by industry (Source: Recorded Future)


The point is that nuance in the data requires corresponding nuance in how you communicate with management, how you tell the story of your security program, and how you present plans for minimizing the risk to your organization.

In the next chapters, I describe five primary risk impacts: legal or compliance failures, operational disruption, brand impairment, financial fraud, and competitive disadvantage.

Afterward, I will walk you through how to measure and communicate risk to executives, connect intelligence to risk through repeatable workflows, use cyber intelligence as an operational security control and a force multiplier for every security function, produce great intelligence (how the sausage is made), and finally, maximize the value of security spending.

Chapter 2

Risk Impacts




After studying cyber risk for the past 20 years and having conversations with hundreds of CISOs in every possible industry, I believe all cyber threats, left unaddressed, lead to five types of business impact:

1. Legal or compliance failure
2. Operational disruption
3. Brand impairment (sometimes referred to as “reputational risk”)
4. Financial fraud
5. Competitive disadvantage

These five categories are important because they reduce uncertainty and increase executive engagement. They make it easier for security leaders and executives to answer “so what?” and “now what?” questions when consuming intelligence and managing cybersecurity programs.

How common are events that fall into these categories of risk impact? The chart below summarizes loss events analyzed in the study of 400 public cybersecurity failures mentioned at the end of the last chapter.




Percentage of loss events, by risk category (Source: Recorded Future)

Note, however, that these frequencies vary tremendously across industries. Because of different likelihoods of occurrence and different impacts, every organization will prioritize risk types and mitigation efforts its own way.

Industry sector typically plays the greatest role in prioritization. For example, financial services companies spend considerably on controls to manage financial fraud. They are also concerned with legal and compliance failures (such as the theft of PII), brand impairment, and operational disruption such as that caused by ransomware. They are generally much less worried about cyber campaigns that could result in a competitive disadvantage from lost intellectual property.

It is crucial to understand a business and its sensitivity to different risk impacts in order to optimize its security programs. For instance, a security team designing intelligence requirements (sometimes referred to as “priority intelligence requirements” or PIRs) must understand relevant risk impacts before working backward to the intelligence that can mitigate those risks.



A security team worried about a particular risk impact can work backward to the related TTPs

Risk Impacts and Adversary TTPs

When a security team is worried about a particular risk impact, it can work backward to the tools and tactics (we say “TTPs,” or “tactics, techniques, and procedures” in industry jargon) that adversaries employ to damage systems and networks, leading to risk impacts. That analysis points to specific security controls and intelligence requirements that can be used to thwart the TTPs most dangerous to the organization.

An Example: Acme Financial Services

Let’s work with a simple example: Acme Financial Services. Acme has \$50 billion under management and provides retail services, investment management, and wholesale services. Acme’s leaders realize the time has come to improve security through better risk assessments and relevant intelligence. Where does the company start?

To continue international operations and avoid fines and additional oversight, Acme spends significant resources on governance and compliance. Also, Acme’s reputation is a priority in an industry with plenty of competition. Finally, Acme needs to avoid operational disruption so it can trade and service its clients reliably.

Of course, Acme is also concerned with financial loss from fraud. But because Acme has spent significant resources build-

ing controls to minimize financial fraud, it is a distant fourth in the company's risk impact prioritization. Competitive disadvantage comes in last because the company isn't competing in markets in which government-sponsored adversaries or competitors are likely to steal intellectual property and trade secrets.

This risk identification is remarkable because Acme's CISO, Nina, performed it without needing to meet with each line of business head. Her background knowledge of Acme's business enabled her to efficiently identify the risks that matter to the firm.


Now, Nina is focused on mitigating risk impacts from legal or compliance failure, brand impairment, and operational disruption (in that order). Her security team works backward to identify attack types with high impact or proximity (meaning occurrence at competitors and other organizations similar to Acme). As fortune would have it, Recorded Future's intelligence cloud is available to help Acme identify the adversaries and associated TTPs most likely to cause disruptions.

The figure on the next page shows simple examples of how Nina and her team were able to pinpoint the TTPs most likely to have a high impact on Acme, together with the threat actors most likely to develop and use them. This intelligence will allow Acme to prioritize security investments effectively and get the biggest "bang for its buck" in reducing risk.

Reducing Uncertainty with Intelligence

Good executives of all types appreciate how *business* intelligence can help them create a durable decision advantage over rivals and, by extension, a competitive advantage. But few are aware of how *threat* intelligence can produce equally important decision advantages.

In my view, "uncertainty" is the 2023 word of the year. Uncertainty is multiplying because of wars, conflicts, inflation, the results of climate change, high-altitude balloons, natural disasters, political strife, evolving legislation, macro-economic downturns, and threat actors using generative artificial intelligence (AI) in new and harmful ways.



Nina and her team pinpointed the most important threat actors and their TTPs based on attack types with the highest potential impact on ACME and the greatest proximity (occurrence at similar organizations)

Similarly, cyber and physical threats change daily, sometimes by the hour. Even the U.S. National Security Agency (NSA) struggles to manage the mind-boggling amount of information that needs to be processed. To avoid ocean-boiling, intelligence and security professionals must begin prioritization with one or more of the five risk impacts and work back to security controls to mitigate them.

These business risks don't exist in isolation. They blend in unique ways.

For example, a vehicle manufacturer might be concerned about the threat of nation-state adversaries, particularly China, stealing trade secrets or intellectual property in the form of proprietary code used in electronic vehicle (EV) platforms. Theft of that code could hand over competitive advantage to another EV manufacturer, erasing billions of dollars of future revenue for the unfortunate victim.



Or the theft of the same source code by a financially motivated actor could lead to an online criminal auction. In this scenario, news of the “dark web” auction could do considerable damage to the victim’s brand and reputation with investors, customers, suppliers, government agencies, automotive and technology partners, and others.

Finally, physical and cyber threats can interact. The vehicle manufacturer should conduct an “all hazards assessment” to consider how cyberattacks, terrorism, natural disasters, and ideological conflicts can exacerbate each other and ultimately magnify one or more of the five risk impacts that create organizational loss.

Robust intelligence is often the best way, and sometimes the only way, to reduce uncertainty in situations like these.

For example, mitigating brand impairment requires broad visibility into criminal communities. These communities, which span the web, internet relay chat (IRC), Telegram, and other types of mobile chat apps, involve communication in numer-

ous languages. Intelligence can ensure awareness of the TTPs they use to steal sensitive data, as well as unauthorized data sales and auctions. However, data collection and intelligence analysis need to be broad and timely to produce the desired outcomes.

In summary, intelligence is a very powerful risk management tool when properly instrumented. It can function as an operational control when automated, and as a strategic advantage when put in the hands of capable operators and analysts. The key is to ensure that the right intelligence, built on the right data, is available to the right people to combat the most pressing risk impacts.

In the third section of this book (chapters 11-14) we will talk about how to design and manage an intelligence program that improves security and reduces risk. But now, let's talk more about the five types of risk impacts.

Chapter 3

Legal or Compliance Failure



Of the five risk impacts I discuss in this book, the “Legal or Compliance Failure” category is most familiar to a board of directors. Board members often have experience or training related to legal and compliance issues, and are usually well aware of examples of corporate malfeasance that led to legislation like the Sarbanes Oxley Act (SOX).⁵ They understand accounting scandals and recognize that poor governance breeds lawsuits and potentially jail time.⁶

Security compliance is still unfamiliar terrain for many boards, particularly when their organization operates in multiple geographies. However, board members understand that this risk category is important because of the detrimental effects generated by failures, especially fines and costs for additional oversight.

5. https://www.law.cornell.edu/wex/sarbanes-oxley_act

6. <https://www.sec.gov/news/press/2006/2006-58.htm>

The U.S. Security and Exchange Commission's proposed rules for cybersecurity risk management are a step toward improved security and governance. An SEC press release summarized some of these new rules:

The proposed amendments would require, among other things, current reporting about material cybersecurity incidents and periodic reporting to provide updates about previously reported cybersecurity incidents. The proposal also would require periodic reporting about a registrant's policies and procedures to identify and manage cybersecurity risks; the registrant's board of directors' oversight of cybersecurity risk; and management's role and expertise in assessing and managing cybersecurity risk and implementing cybersecurity policies and procedures. The proposal further would require annual reporting or certain proxy disclosure about the board of directors' cybersecurity expertise, if any.⁷

When the SEC (and other regulatory agencies around the world) hold board members accountable for managing security risks and implementing cybersecurity policies, you can bet their interest in preventing compliance failures will grow. Also, there will be more pressure on IT security leaders to communicate with senior executives and boards of directors using messages that are simple but convey much more cyber nuance.

More Regulations, More Failures, Rising Costs of Non-Compliance

A brief tour around the world of recent legal and compliance failures illustrates their ubiquity and the rising costs of non-compliance. It also underscores the importance of threat intelligence to properly assess and manage cyber risk, which is now, more than ever, business risk.

2018

- Ticketmaster was fined \$1.25 million under Article 5 of the GDPR and found negligent due to a chatbot implementation on its payments page that resulted in the compromise of 60,000 payment card details and 1,000 financial loss victims.⁸

7. <https://www.sec.gov/news/press-release/2022-39>

8. <https://cs.brown.edu/courses/csci2390/2021/assign/gdpr/mzhan104-knelson9-zlee8-ticketmaster.pdf>

2019

- French data privacy authority CNIL fined Bouygues (a telecom company) 250,000 euros “for a security breach that affected the personal data of around two million clients for over two years.”⁹

2021

- Nigeria’s National Information Technology Development Agency (NITDA) fined Electronic Settlement Limited (ESL) 5M Naira under the Nigeria Data Protection Regulations (NDPR), and increased IT oversight for six months to monitor new controls including a “clear data security and governance document” for ESL vendors and suppliers.¹⁰
- The Dutch Supervisory Authority (DSA) fined Transavia (a French airline) for poor security after a hacker downloaded the PII on 83,000 individuals. DSA faulted Transavia for a lack of multi-factor authentication (MFA) and a lack of network segmentation (now fashionably referred to as “zero trust”).¹¹

2022

- Optus (an Australian telecom company) sustained a data breach that “prompted lawmakers to introduce the Privacy Legislation Amendment Bill 2022, which increases fines to AU\$50 million when companies sustain repeated data breaches.”¹²
- Singapore’s Personal Data Protection Commission (PDPC) fined RedMart (an online supermarket) 72,000 Singapore dollars (about \$54,000 USD at the time) for a violation of Section 24 of the Personal Data Protection Act 2012. The PDPC found RedMart failed to implement reasonable security arrangements.

9. <https://www.reuters.com/article/us-france-bouygues-fine/french-watchdog-fines-bouygues-for-data-security-breach-idUSKCN1OQ0Q4>

10. <https://www.nigeriacommunicationsweek.com.ng/nitda-slams-n5m-on-esl-over-data-breach/>

11. https://edpb.europa.eu/news/national-news/2021/dutch-sa-fines-transavia-poor-personal-data-security_en

12. <https://iapp.org/news/a/a-look-back-at-privacy-and-data-protection-in-2022/>

Specifically, during an IT migration project, the company failed to encrypt its customer database or implement password authentication to it.¹³

2023

- DISH, the satellite broadcast company, was the victim of a ransomware attack that also resulted in a data breach. The company told regulators that the data of 296,851 people was affected by the incident. In breach notification letters they confirmed that personal data was involved, including driver's license numbers. As a result, DISH faced a class-action lawsuit.¹⁴
- EyeMed Vision Care, an eye insurance provider, paid a \$2.5 million fine due to a 2020 data breach that involved the personally identifiable information of 2.1 million people. The 2023 payment added to a total of almost \$8 million for violations of HIPAA and various state regulations.¹⁵

Data Privacy: The New Legal Frontier

In 2023, global regulatory regimes are accelerating data privacy protections. Argentina is revisiting its outdated Personal Data Protection Act (Ley de Protección de los Datos Personales). Canada is legislating the Consumer Privacy Protection Act, the Personal Information and Data Protection Tribunal Act, and the Artificial Intelligence and Data Act. India and Indonesia both proposed digital personal data protection bills. Following the lead of California, Virginia, and Colorado's, Connecticut and Utah passed data privacy legislation in 2022.

Security and privacy legislation seek different outcomes, but they are closely linked. Regulatory authorities are increasingly mandating and scrutinizing security practices to accomplish data privacy.

13. <https://www.dataguidance.com/news/singapore-pdpc-fines-redmart-sgd-72000-failure-ensure>

14. <https://therecord.media/people-affected-by-dish-data-breach>, <https://rosenlegal.com/wp-content/uploads/2023/03/DISH-Complaint-Web-4856-6712-9433-v.1.pdf>

15. <https://therecord.media/eyemed-data-breach-settlement-four-states>

In the United States, while Congress debates a national data privacy law, progress is emerging through a patchwork of new state legislation and federal agencies with the power to enforce compliance. The SEC wields the big stick in financial services. The U.S. Federal Trade Commission (FTC) has become a de facto enforcer of corporate data privacy because of its historical role monitoring compliance with legislation on advertising.

Critical industries, as designated by the U.S. federal government, also have a patchwork of mandates for compliance. Electric utilities, for example, are regulated by the Federal Energy Regulatory Commission (FERC), which enforces compliance with standards defined by the North American Electric Reliability Corporation (NERC).¹⁶

Compliance frameworks range from guides to best practices like the NIST CSF to mandates like those in HIPAA for health-care and PCI DSS for vendors that process payment cards.

Legal Failures

Legal failures and compliance failures are not always synonymous. While a failure in one category may imply a failure in the other, they can be separate.

For example, the U.S. Treasury's Office of Foreign Asset Control (OFAC) regularly sanctions cybercrime groups like Trickbot because they seek to attack critical infrastructure, exploit the international financial system, or support terrorism.¹⁷ When an organization with a low tolerance for operational disruption is infected with ransomware, executives may decide to pay the ransom. But paying a ransomware group that has been sanctioned is illegal. In the fog of an emergency, engaging with a seemingly anonymous actor online and remitting cryptocurrency can feel benign, but that payment may constitute a legal failure.

Legal compliance can be unexpectedly difficult. Consider a legal requirement from Saudi Central Bank (formerly the Saudi Arabian Monetary Authority, or SAMA), the financial

16. <https://www.ferc.gov/industries-data/electric/industry-activities/nerc-standards>, <https://www.ferc.gov/industries-data/electric/industry-activities/cyber-and-grid-security>

17. <https://home.treasury.gov/news/press-releases/jy1256>

regulator in Saudi Arabia. It mandates, via its cybersecurity framework and the National Cybersecurity Authority's Essential Cybersecurity Controls, that Saudi banks rapidly eliminate cyber properties (mobile apps, websites, and so on) that impersonate a Saudi bank brand.¹⁸ However, the internet's plumbing often doesn't facilitate impersonation takedown requests very well. If a domain is registered with a Chinese registrar, for example, there is a process to request to remove the offending domain, but it can take days to evaluate and act on the request.

How Intelligence Can Prevent Legal and Compliance Failures

Great, you say. I understand there are significant legal and compliance cyber mandates. Our company has a whole governance, risk, and compliance (GRC) group to ensure we are meeting expectations. Why is intelligence required to mitigate this risk?

In the summer of 2018, I would wager that executives at British Airways were confident that the airline was GDPR compliant. They understood the regulation and had deployed security controls to meet its requirements. Yet British Airways' website was compromised for approximately two weeks. A few lines of inserted code skimmed the details of about 500,000 payment cards as customers booked flights. The result was a GDPR fine of millions of pounds.

How could this happen? The issue is that the GDPR standard is largely static, with only minor annual updates, while adversary attacks evolve rapidly, making initially deployed security controls outdated and ineffective. In the case of British Airways, a clever group of cyber thieves developed and deployed a toolkit called "Magecart" that the airline's "compliant" defenses could not detect.

Recorded Future happens to track Magecart victims by identifying Magecart code implants in their websites. Understanding the full Magecart attack life cycle is important intelligence for an organization operating an e-commerce website. If British Airways had known about the tactics

18. <https://nca.gov.sa/ecc-en.pdf>

Magecart gangsters use to gain unauthorized access to websites and embed their code, they might have reduced the risk of running afoul of the GDPR.

An organization that prioritizes mitigation of legal or compliance failure risk can start by building a list of applicable laws and compliance mandates. Next, it can identify adversary TTPs capable of causing legal and compliance failures. It can commence TTP monitoring and alerting once its security team is confident about their processes for sourcing, collecting, and analyzing data.

In a later chapter I will discuss the concept of “relevant threat deltas” (RTDs). Once a threat type or category is identified, tracking movement in that category is the work that matters. Sometimes the threat movement is imperceptible, and sometimes it’s massive.

Of course, there is also an issue of alignment and communication between security and business leaders. A CISO understands the nuances of intelligence and security control implementations. Members of the board want binary risk statements (“This thing is a risk, this other thing is not a risk.”) To avoid misunderstandings, CISOs must learn to quantify risks and talk the language of business. I will discuss that in the second section of this book.

AI Advances and Cybercrime Advantages

For adversaries, using an .hta file instead of a .pdf file to deliver malicious code via email is a relatively minor change in tactics.

ChatGPT, as a generative artificial intelligence (AI) chatbot, is an enormous leap forward for adversaries that craft duplicitous messages for phishing, vishing, smishing, and other social engineering attacks. It is particularly helpful for adversaries who don’t speak the language of their target well.

Ironically, ransomware actors are now referencing data privacy governance to cajole victims into paying ransom. A blog post from cybersecurity firm Redacted describes the BianLian ransomware gang’s latest ploy:

In several instances, BianLian made reference to legal and regulatory issues a victim would face were it to become public that the organization had suffered a breach. The group has also gone so far as to include specific references to the subsections of several laws and statutes. While the applicability of the laws (to the victim and their data) referenced by BianLian would need to be assessed by the courts, at first glance, the laws referenced by the actors did in fact correspond to the jurisdiction where the victim was located. This attention to detail shows that the criminal gang is taking the extra time to tailor threats to their victims to maximize the pressure to pay the ransom.¹⁹

Most regulatory mandates are aimed at data confidentiality. All methods of accomplishing unauthorized network access are a threat.

Defenders need to anticipate what will happen when an attacker gains unauthorized access. How long will it take to detect the trespasser? What's the time to remediation? They also need to foresee problems that can be caused by different classes of attackers: insiders, extortionists employing ransomware to steal data, nation-state adversaries exfiltrating data for foreign government use, hacktivists looking to embarrass the victim, and financially motivated actors who sell stolen data to the highest bidder in digital markets.



19. <https://redacted.com/blog/bianlian-ransomware-gang-continues-to-evolve/>

Among these actor groups, the financially motivated are likely the largest threat to businesses (although precise statistics are elusive). A sample of criminal market advertisements captured by Recorded Future in Q1 2023 reveal international opportunities in every industry vertical:

- Kernelware, a member of mid-tier BreachForums, was selling 160 GB of data from a well-known Taiwanese multinational hardware and electronics company.
- Ronyking247, a member of the top-tier forum Exploit, was selling a database containing PII of employees from several unspecified US and UK companies.
- Fullcrypt, a member of the top-tier forum Ramp, sold unauthorized access to the network of an unspecified Israeli company.
- 740182, a member of the Chinese-language Exchange Market, was selling a database containing eight records of Microsoft offline store managers based in China.
- El84, a member of the top-tier forum XSS, was selling domain administrator access to the Brazilian branch of a well-known Japanese auto manufacturer.
- TwoFactor, a member of the mid-tier BreachForums, was selling access to the email account of a high-ranking police officer in India.
- MODEGYPT, also a member of BreachForums, was selling 2 million records from a database allegedly belonging to the Egyptian Ministry of Health.
- Xinploiter, a member of Exploit, was selling unrestricted access to a server owned by an unspecified law firm in the United Arab Emirates.
- kali88, a member of the Chinese-language marketplace Chang'An Sleepless Night, was selling four sets of PII from organizations in China, Vietnam, and Taiwan.

How can a business respond to and learn from control failures if it lacks even the basic intelligence necessary to identify when its own or its supplier's data is being auctioned? A

comprehensive data sourcing and collection effort that spans everything from members-only Russian-language digital clubs to English-, Chinese-, Arabic-, and Hebrew-language Telegram channels is required.

How can a business understand the risk of a legal or compliance failure if it does not have daily insight into how adversary tactics are faring against its security program?

In a 2019 interview with *The Wall Street Journal*, U.K. information commissioner Elizabeth Denham pointed to multiple variables when calculating GDPR fines, specifically “a company’s size, the number of people affected, and the length of time that hackers had access to data before they were detected.”²⁰


Businesses need to invest in intelligence because it acts as a security control for breach detection. The availability of intelligence is particularly relevant as regulators judge security in hindsight. Denham went on to say, “Our focus is whether or not there was adequate, reasonable, consistent, effective data security to protect people’s data.” Those are subjective labels, which many executives translate as an industry litmus test: “Are our competitors investing in this control?”

Lacking intelligence, “adequate, reasonable, consistent, effective security” is difficult to assess. Similarly, regulatory fines and class action lawsuits for personal data breaches contain an element of timing. After a breach is reported, financial losses may not be incurred for years, at which point they may become the responsibility of a new leadership team.

20. <https://www.wsj.com/articles/u-k-regulator-on-why-it-is-pursuing-record-fines-against-ba-marriott-11562751006>

Chapter 4

Operational Disruption



Operational disruptions are unexpected events that prevent an organization from transacting business with customers or suppliers, building and transporting products, promoting and advertising goods and services, developing new offerings, calculating revenues and expenses, or carrying out other business activities. I am going to focus on cyber threats here, but physical events such as earthquakes, fires, floods, hurricanes, protests, and pandemics can also cause operational disruptions with similar results for victims.

Software as a service (SaaS) and cloud companies lose customers when their operations are impaired. Manufacturers are generally intolerant of downtime. Disrupted access to medical records at healthcare organizations can have life-threatening consequences. Vital government services can come to a halt. In these and other verticals, when the phones start ringing, the pressure builds quickly, because the impact on revenue, customer goodwill, and safety can start in seconds.

In fact, operational disruptions can have far-reaching effects and cause more types of damage than most people can imagine. A 2018 N-1A statement from the Dreyfus Corporation provides a long (and probably not 100% complete) list of the impacts cyberattacks can have on a mutual fund company:

Cybersecurity incidents affecting the Manager, Subadviser(s), Transfer Agent or Custodian or other service providers such as financial intermediaries have the ability to cause disruptions and impact business operations, potentially resulting in financial losses, including by interference with a fund's ability to calculate its NAV [net asset value]; impediments to trading for a fund's portfolio; the inability of fund shareholders to transact business with the fund; violations of applicable privacy, data security or other laws; regulatory fines and penalties; reputational damage; reimbursement or other compensation or remediation costs; legal fees; or additional compliance costs. Similar adverse consequences could result from cybersecurity incidents affecting issuers of securities in which a fund invests, counterparties with which the fund engages in transactions, governmental and other regulatory authorities, exchange and other financial market operators, banks, brokers, dealers, insurance companies and other financial institutions and other parties.²¹

The last part of this excerpt highlights the fact that organizations are also vulnerable to attacks on their suppliers and business partners. When operations are disrupted, the damage is rarely confined to the initial target. In February 2023, Applied Materials made headlines when the technology company announced a \$250 million hit to quarterly sales due to a cyberattack on a key supplier. These third- and fourth-party exposures are difficult to remediate. As one CISO put it, "Tell me when a supplier is compromised so I can immediately cut access and determine how best to help them."²²

Operational disruption isn't always a self-contained risk impact. A successful cyberattack can begin with operational disruption and end in a compliance failure and associated reg-

21. <https://www.sec.gov/Archives/edgar/data/819940/000081994018000017/lp1d1fi.htm#161>

22. <https://www.bizjournals.com/austin/news/2023/02/19/applied-materials-250m-revenue-hit-supplier-hack.html>

ulatory fines. The Wall Street Journal documented the fines levied on Centric Health Ltd., Discord Inc., and Interserve Group Ltd., when alleged missteps in response and recovery activities after an incident led to the inadvertent (and illegal) destruction of patient data.²³

Most organizations need to be alert to two distinct types of operational disruption.

Disrupting Networks and Services

Threats such as distributed denial-of-service (DDoS) attacks overwhelm networks, web, DNS, and database servers, and web applications. A wide variety of flooding, amplification, and reflection attacks can overwhelm networks, while application-level attacks exhaust application services (for example, by submitting queries that tie up massive amounts of computing resources).

Naturally, a little creativity goes a long way in accomplishing operational disruption. Website defacements are generally regarded as a nuisance,²⁴ but taking a web application offline via DDoS attacks is potentially disastrous for B2C firms, particularly in transportation, and for B2B infrastructure including cloud and SaaS platforms.²⁵

Destroying Data and Disabling Systems

Operational disruptions can also be caused by attacks that encrypt or corrupt data or disable computing systems. The most prominent example today is ransomware that encrypts files until a ransom is paid (or sometimes forever, despite a ransom being paid). Over the past decade, ransomware has gained significant attention from the mainstream media because the adverse effects on victims can be pronounced and serious.

Since the first version of this book was published in 2020,

23. <https://www.wsj.com/articles/inadvertent-data-destruction-after-a-cyberattack-can-violate-eu-privacy-rules-a796d8e>, <https://www.nist.gov/cyberframework/online-learning/five-functions>

24. <https://therecord.media/danish-hospitals-hit-by-cyberattack-from-anonymous-sudan>

25. <https://therecord.media/cloudflare-says-it-stopped-largest-ddos-attack-on-record>

ransomware has been on a tear. The Colonial Pipeline attack²⁶ was a watershed event in that diplomatic pressure may have triggered a move away from targeting high-profile U.S. enterprises and toward companies in other parts of the world and government agencies in small countries.²⁷

In 2023, Capita, a British business services outsourcing company, said that remediating a recent ransomware attack it was victim to would cost \$25 million, specifically for “specialist professional fees, recovery and remediation costs and investment to reinforce Capita’s cyber security environment.”²⁸

But ransomware is not the only threat that can cause long-term or permanent damage. Malware can be used to disable network, database, and application servers, and wipe the content of databases and file servers. Moreover, malicious software can change settings in ways that go beyond disruption to catastrophic failure, for example, by causing centrifuges to tear themselves apart violently, or shutting down safety systems in a petrochemical plant.

Threats to Operations Are Evolving Rapidly

While DDoS and ransomware attacks have been around for a long time, the techniques are not static. On the contrary, threat actors are continually deploying new variations that defeat existing controls.

New classes of DDoS attacks have emerged recently. For example, inventory exhaustion and hoarding attacks temporarily fill shopping carts so e-commerce companies falsely appear to be sold out of popular items.

Now, in Q1 2023, ransomware payments from victims are declining.²⁹ You might think this is good news, and mostly

26. <https://www.bloomberg.com/news/articles/2021-06-04/hackers-breached-colonial-pipeline-using-compromised-password?sref=D1QuWRle>

27. <https://therecord.media/spanish-amusement-park-giant-hit-with-cyberattack>, <https://therecord.media/medibank-says-it-will-not-pay-ransom-in-hack-that-impacted-9-7-million-customers>, <https://therecord.media/ransomware-gang-threatens-to-overthrow-new-costa-rica-government-raises-demand-to-20-million>

28. <https://therecord.media/capita-ransomware-incident-response-cost>

29. <https://www.wsj.com/articles/ransomware-attacks-decline-as-new-defenses-countermeasures-thwart-hackers-23b918a3>

it is. However, it may force threat actors to innovate once again. Just as global adoption of EMV chips in payment cards reduced the monetization opportunities for stolen card data (referred to in criminal parlance as “dumps”),³⁰ an improved playbook for remediating ransomware may trigger a more sinister approach from adversaries. At the risk of incorrect prognostication, what comes next may involve code that disables physical systems by way of firmware, such as master boot record (MBR) destruction or low-level firmware flashing.

In addition, the spectacular growth in the number of internet-of-things (IoT) devices will open many opportunities for adversaries to disrupt operations, not only in factories, utilities, and transportation networks, but also in cities, hospitals, and even in vehicles and homes.

Now, all is not doom and gloom. Only a subset of these threats to operations are relevant to any one organization. You can use intelligence and a risk-based approach to security to determine which threats are actual risks to your operations, and further, which risks are large enough to justify an investment in controls and staffing.

30. <https://www.darkreading.com/risk/ransomware-carding-s-replacement-for-the-criminal-masses>

Chapter 5

Brand Impairment



could use “reputational risk” here, but “impairment” is a better term in this case. It suggests a malfunction that leaves an operation running but hobbled.

Brand cultivation has never been more important than in these days of social media sniping and high-velocity news. One misstep, one exposure of a controversial idea from a private conversation, or one malicious actor impersonating an executive can wreak havoc or antagonize half the population. Cybersecurity is now a material piece of managing a brand.

This is particularly true in the age of environmental, social, and governance (ESG) strategy. For many stakeholders, good security is an aspect of responsible behavior. Employees, vendors, and customers want to work for and do business with organizations that invest in security to ensure that data breaches don’t transpire and that the personal data and privacy of stakeholders are well supported.

In the early days of the internet, positive social behavior involved investing in security to prevent threat actors from using corporate assets to launch DDoS attacks that victimized other organizations or slowed the internet. Today, positive social behavior includes securing products and services that are part of larger supply chains. It also means constructing sufficient identity verification systems so that unauthorized access is no longer possible via common methods (stolen credentials, for example).

Good governance means asking intelligent oversight questions, particularly about the controls that are safeguarding customer and employee PII. Allowing PII to fall into the wrong hands is a fast-track to brand impairment.

Another express lane to brand impairment is public data extortion. In the new breed of ransomware attacks, cyber-criminals not only encrypt files, they also export copies and threaten to expose the data. While audiences tend toward sympathy with the victims of ransomware attacks, the public release of sensitive information is always going to be detrimental to a brand. The increasing danger of these attacks is indicated by the proliferation of ransomware extortion websites (operated behind Tor) such as Karakurt Group Leaks, LockBit 3.0 Leaked Data, Ransomexx, BianLian, Onyx News, Vice Society, Clop Leaks, RandomHouse, and CryptOn.



Ransomware extortion websites like Vice Society impair brand image by increasing the visibility of successful data breaches

Any event that hurts a brand's reputation — whether it's a cyberattack resulting in leaked data, the public exposure of a security or compliance failure, or a business VIP's controversial political stance that quickly populates headlines — can have cascading risk impacts, including further brand-impairing events.

For example, Hindenburg's March 2023 report on Block (provider of Cash App) highlighted alleged compliance violations that harmed the corporation's stock price and ultimately Block's reputation.³¹ It takes a long time to recover reputation and investor confidence.

Uber experienced brand impairment from a cyber incident. In 2018 an employee stole sensitive company data before departing. The information included customer data, driver data, and proprietary technology, all of which was carried off the premises on the employee's personal computer.

In 2020, notable Twitter accounts, including those of Elon Musk, Bill Gates, Jeff Bezos, and Barack Obama, were used to

31. <https://www.bloomberg.com/news/articles/2023-03-23/block-shares-fall-after-hindenburg-says-it-s-short-the-stock?sref=D1QuWRle>

promote a cryptocurrency scam. A spear phishing campaign that purported to originate from IT staff members allowed attackers to obtain employee credentials that enabled them to misuse Twitter’s internal systems.

These incidents exposed gaps in Uber’s and Twitter’s security programs, damaging the companies’ brand equity with employees, customers, and shareholders.

Brands can also be damaged by attacks that never actually touch an organization’s own websites or social media accounts. Common examples are websites that appear to belong to well-known entities but are controlled by threat actors (typosquatting), fraudulent online stores appearing to belong to famous brands, counterfeit apps in app stores, and social media accounts mimicking those of the victim and its executives. These might deliver counterfeit goods (or no goods at all), disseminate disinformation or offensive opinions, or capture customer information to facilitate fraud. Since they are not visible from the organization’s systems, intelligence may be the only way to detect them before they cause serious brand impairment.

**www.converse.com.br.x-bitbucket-pr-759-minuzja-ijw2b
cnplawie.us-5.magentosite.cloud**

References	2
First Reference	May 2, 2023
Latest Reference	May 2, 2023
🕒 New Domain	Added May 2, 2023
Recorded Future Community	Domain 🔗



30
SUSPICIOUS
RISK SCORE

2 of 53 Risk Rules Triggered

[Show recent events](#) or [cyber events](#)

An example of a typosquatted domain, in this case spoofing Converse[.]com, shown in a Recorded Future Intelligence Card

Finally, corporate ideological positions and alliances have new significance when evaluating stakeholder impact, including the risk of future cyberattack victimization.

OpenAI’s ChatGPT provided the following summary of a Coinbase blog post:

In September 2020, Coinbase CEO Brian Armstrong published a blog post outlining the company's mission and values, which included a strong focus on maintaining an apolitical work environment. In the blog post, Armstrong argued that the company should focus on its core mission of creating an open financial system for the world and avoid engaging in broader social and political issues that are not directly related to its business.

This stance generated a significant amount of controversy both internally and externally. Some Coinbase employees felt that the company's position was dismissive of important social and political issues that directly or indirectly affect them and their communities. As a result, they argued that the company should take a more active role in addressing these concerns.

In response to the concerns raised by employees and the broader public, Coinbase offered a severance package to any employee who disagreed with the company's stance and wished to leave. This severance package included four months of salary and six months of health insurance coverage. As a result of this offer, approximately 60 employees (about 5% of the workforce at the time) chose to leave the company.

Shopify experienced a similar situation. The summary follows:

In February 2017, Shopify faced controversy and internal debate over its ideological position when the company decided to continue hosting the online store of Breitbart News, a far-right news organization known for its controversial content. The decision was based on Shopify's commitment to supporting free speech and enabling commerce for a wide range of businesses, even those with differing political views.


This decision generated criticism from some employees, customers, and members of the general public, who argued that Shopify should not support an organization like Breitbart, which they believed promoted divisive and harmful ideas. In response to these concerns, Shopify CEO Tobi Lütke published an open letter explaining the company's stance on the matter.

In the letter, Lütke stated that Shopify's mission was to enable commerce for everyone, even if their views differed from those of the company or its employees. He argued that by supporting a broad range of businesses, Shopify was upholding the principles of free speech and democracy.

Ideology and Cyber Targeting

Ideological positions can have consequences in the modern business world. One of the newer consequences involves cyber targeting.

Hactivists target organizations based on their ideology. These attacks usually involve defacing websites or stealing and disclosing data damaging to the organization's reputation. In the first quarter of 2023 alone, such attacks have been attributed to more than 40 threat actor groups, including Anonymous, Killnet, KelvinSecTeam, Ashiyane Digital Security, KillMilk, CtrlSec, Cyber Partisans, Edalat-e Ali, and 1877 Team.



An example of hacktivists publishing stolen information from a religious organization.

Nation-states as well as hacktivists are targeting organizations based on their alliances or publicly disclosed social positions. North Korea caused widespread destruction at Sony via a cyber wiper over a perceived slight to “Supreme Leader” Kim Jong Un.³² Iran wreaked havoc inside the Sands Corporation after its owner, Sheldon Adelson, made perceived offensive comments.³³

China is targeting international organizations based on any alliance or policy position that is or may become relevant to the Chinese Communist Party. In the past, companies assumed they would fall in the crosshairs of the People’s Liberation Army (PLA) or the Ministry of State Security (MSS) only if they operated in an industry important to China’s

32. <https://www.justice.gov/opa/pr/north-korean-regime-backed-programmer-charged-conspiracy-conduct-multiple-cyber-attacks-and>

33. <https://www.bloomberg.com/news/articles/2014-12-11/iranian-hackers-hit-sheldon-adelsons-sands-casino-in-las-vegas>

Five-Year Plan.³⁴ However, China's intelligence resources are enormous, and today no potential information advantage is beyond their scope.³⁵

And I mean *no* potential information advantage, even some that are impossible for us to divine. In Q1 2023, Insikt Group, Recorded Future's research arm, reported on a Chinese government-sponsored group called Threat Actor Group 22 (TAG22) that had compromised infrastructure belonging to the Palestinian Ministry of Awqaf and Religious Affairs. In Q4 2022 Insikt Group reported on RedDelta, another Chinese state-sponsored group, which compromised the network of Spain's Ministry of Tourism and Trade. One theory is that the Chinese government values intelligence about any policy position by a foreign government agency related to issues such as religious evangelism and Covid.

The same is true for Russia. There are multiple examples of Russian cybercriminals moving off the sidelines to directly support Russia in its war with Ukraine.³⁶ Companies cannot afford to underestimate the risk to brands inherent in taking positions such as supporting Ukraine in this conflict. The risk may be acceptable, but organizations should weigh the potential costs and consider options for mitigation.

Sometimes disclosing cyber risks can help protect your brand. For example, if you are an executive at a publicly traded company that takes stances on issues or announces policies that may be controversial, you may want to advise shareholders that those public positions could invite increased cyberattacks. In the future, regulators may require such public disclosure by publicly traded companies.


34. <https://www.globalpolicywatch.com/2021/04/chinas-14th-five-year-plan-2021-2025-signposts-for-doing-business-in-china/>

35. <https://www.wsj.com/articles/rise-of-open-source-intelligence-tests-u-s-spies-11670710806>

36. <https://www.recordedfuture.com/dark-covenant-2-cybercrime-russian-state-war-ukraine>

Chapter 6

Financial Fraud



Financial fraud has always been an unfortunate tax on society. Relatively old-fashioned check and payment card thefts still plague financial services companies and their customers. More recently, in the process of bringing valuable new products to market, the fintech industry has created new fraud monetization opportunities. The global adoption of payment cards with EMV chips has reduced the frequency of traditional “carding” schemes,³⁷ but money transfer services like Zelle, Venmo, Wire, Chime, and CashApp and merchant services like Stripe and Square, allow fraudsters to target unsuspecting customers with social engineering scams, and data dumps containing the magnetic stripe information from credit cards are commonly sold on cybercriminal marketplaces.

Additionally, cryptocurrency exchanges are under constant attack.³⁸ The attackers include both rank-and-file financially

37. <https://www.nfcw.com/2022/06/13/377453/more-than-90-of-card-present-payments-worldwide-were-made-using-emv-chip-cards-in-2021/>

38. <https://cointelegraph.com/explained/the-biggest-crypto-heists-of-all-time>

motivated criminals and nation-state actors. North Korea in particular has been actively victimizing exchanges to generate illicit profits for the regime.³⁹ In early 2023, for example, Recorded Future's Insikt Group observed DPRK-sponsored groups targeting financial services companies in Southeast Asia and the United States with domain typosquatting and phishing attacks. These companies are engaged in services that range from retail banking to venture capital and private equity.

Financial services companies have their hands full trying to protect their customers, but of course, even the most rigorous technical controls will inevitably fall short of protecting humans from themselves.

Outside of fraud against consumers, business email compromise (BEC) is the quickest adversarial route to a large payday. Stealing directly from accounts payable is relatively straightforward, and the losses may be substantial.

These attacks are even more pernicious when combined with unauthorized access to email accounts. If a BEC attack can compromise administrator accounts via phishing or malware, the attacker can control the internal flow of information between the finance department and everyone else. For example, threat actors have sent emails from a CEO's account that the finance department accepts without question because the sender's address is valid.

Here are examples of historical BEC victims, many of which involve significant financial losses for the victims.

Crelan Bank

In 2016 Crelan, a Belgian bank, suffered a BEC attack that resulted in a loss of €70 million (\$75.8 million). Cybercriminals impersonated bank executives and used social engineering tactics to deceive employees and persuade them to transfer money to fraudulent accounts. Crelan discovered the attack during an internal audit and subsequently increased security measures to protect against future threats.

39. <https://www.amazon.com/Lazarus-Heist-Based-Hit-podcast-ebook/dp/B09QLTPPBW>

Nikkei

In 2019, Japanese media company Nikkei reported a BEC attack that caused a loss of \$29 million. The company revealed that an employee of its U.S. subsidiary had transferred the funds to a fraudulent account after receiving a fake email supposedly from a Nikkei executive. Nikkei worked with law enforcement authorities to investigate the incident and recover the lost funds.

Scoular

In 2016, United States-based agriculture company Scoular lost \$17.2 million in a BEC scam. The attackers impersonated Scoular's CEO and sent an email to a company controller instructing them to wire funds to a bank account in China for a supposed acquisition. The controller followed the instructions, transferring three separate payments to the fraudulent account before realizing it was a scam.

FACC

In 2016, Austrian aerospace manufacturer FACC was hit by a BEC attack that resulted in a loss of €42 million (\$47 million). The attackers impersonated the CEO and sent emails to the finance department, requesting a wire transfer for a fake acquisition project. FACC discovered the attack and disclosed the incident, after which its share price dropped significantly.

A more recent example involves a bank manager in Hong Kong who in 2020 received a deepfake (AI-generated) call. The voice was familiar and he transferred \$35 million to unauthorized parties.⁴⁰

Several Recorded Future clients have related their struggles to contain BEC attempts. The attacks are often focused on foreign subsidiaries whose employees do not speak English as a first language. To capitalize on this language barrier, the attackers use deepfake technology to replicate the voice of the CEO or CFO. A well-timed phone call with a deepfake voice is sometimes enough to induce an employee in finance to initiate an immediate payment to a spurious vendor controlled by a BEC actor.

40. <https://www.forbes.com/sites/thomasbrewster/2021/10/14/huge-bank-fraud-uses-deep-fake-voice-tech-to-steal-millions/>

Artificial intelligence is only going to increase the velocity and precision of BEC attacks. The majority of these attacks have historically originated in West Africa.⁴¹ Artificial intelligence removes previous barriers to grammatically correct English. Additionally, easy-to-produce deepfake videos impersonating CEOs and other executives are on the near horizon. Detecting deepfakes is still a work in progress, and adoption of new technologies that support fraud are almost certain to accelerate. Social engineering is about to undergo a dramatic capability improvement and businesses are largely unprepared to coordinate effective detection and response.

Preventing corporate financial fraud begins with well-designed financial controls, such as requiring two or more individuals to disburse funds or approve and pay invoices. But even organizations with solid controls may find it difficult to apply them consistently after mergers and acquisitions, especially in cases involving foreign entities.

Intelligence also plays a crucial role in detection and remediation. Every organization is probed daily and enterprises are attacked by the second. Technical security controls block the majority of malicious attempts, and the digital exhaust that these controls capture is valuable when correlated with external threat intelligence. These digital traces are particularly important when proactively addressing BEC attacks.

Security professionals love to report on the number of emails blocked by their email security gateway, but often they neglect to mine the gold in the quarantined messages. These offer valuable hints about how threat actors are targeting employees.

Mining email metadata can also improve employee training, a necessity when social engineering attacks powered by AI capabilities are making phishing emails harder to detect. Ultimately, only humans with the right data can prevent well-crafted social engineering attempts.

Organizations typically look at security training as a periodic compliance obligation instead of a valuable opportunity to educate employees with compelling content that adds

41. <https://vdocuments.mx/agari-cyber-intelligence-division.html?page=2>

value, instills confidence, and contributes to a committed workforce.⁴²

Training should cover cybersecurity hygiene at home as well as in the office. For example, lessons learned about the dangers of downloading malicious Minecraft plugins by kids, identifying AI-generated deepfakes on social media, and assessing links contained in SMS messages will carry over into the work environment.

42. <https://www.wsj.com/articles/microsoft-employees-are-hooked-on-the-companys-training-videos-c8684a1>

Chapter 7

Competitive Disadvantage



The Chinese government, working through agencies associated with the PLA and the MSS and via contractors, has been stealing data for at least the past 25 years.⁴³ The resources and scale of China's government-sponsored cyber intrusions, and corresponding data theft, are breathtaking.

The FBI estimates that every year the cost to the U.S. economy of counterfeit goods, pirated software, and theft of trade secrets by China is between \$225 billion and \$600 billion.⁴⁴ Cyber-enabled data heists represent a critical majority of the stolen intellectual property (IP).


China targets government and ministry departments of many countries to achieve an information advantage in negotiations and policy discussions. It targets the private sector to steal data that creates competitive advantages for state-sponsored or subsidized companies, and often for entire industries such as semiconductor manufacturing and quantum computing.

43. <https://www.cfr.org/cyber-operations/titan-rain>

44. <https://www.fbi.gov/file-repository/china-exec-summary-risk-to-corporate-america-2019.pdf>

China isn't the only offender — plenty of countries engage in regular data theft — but China uses its significant offensive capabilities to create an asymmetric advantage. The United States is still a technological leader in most fields, so China has much to gain by accelerating domestic industry development through covert IP transfer from U.S. companies.

China is an equal-opportunity targeter. Virtually every country contains data relevant to some facet of China's domestic growth or geopolitical machinations. For example, in the first quarter of 2023, Insikt Group watched threat activity group RedFoxtro (attributed to Unit 69010 of the PLA's Strategic Support Force Network System Department) target and compromise multiple Indian organizations. These included the state-owned telecommunications provider Bharat Sanchar Nigam Limited (BSNL), the aerospace and defense company Alpha-Elsec, and the District Cooperative Central Bank.



An example of the Diamond Model of Intrusion Analysis in the Recorded Future Intelligence Cloud. It outlines the malicious infrastructure and techniques that RedFoxtro, a Chinese state-sponsored threat activity group, used to attack Indian organizations in various industries in the first quarter of 2023.

PLA and MSS cyber tools and tactics evolve to the extent required to accomplish an objective, but generally even tools with mileage on them are still effective. Routers and switches have historically acted as useful choke points for large-scale data acquisition, and as more computing infrastructure is outsourced, email and database platforms become attractive targets.

Russian hackers have also been detected searching for intellectual property, particularly in aerospace, defense, and energy-related industries.⁴⁵

In these and other attacks that affect competitive advantage, we can see three overlapping motivations:

1. **Military strength.** Governments seek to change the military balance of power in existing and future conflicts by strengthening their defense industrial base (DIB), particularly in the midst of arms races.⁴⁶ Strengthening the DIB can have economic and diplomatic impacts as well as immediate military ones through increased arms sales to client countries and neutrals.
2. **Technology leadership.** Governments and companies strive to achieve technological superiority in emerging fields such as semiconductors, computer and communications equipment, AI, biotechnology, pharmaceutical development, and quantum technology. Looting engineering designs, proprietary manufacturing techniques, and scientific research can speed up market penetration, increase employment, and help new corporations challenge existing industry leaders.
3. **Quick profits.** Some companies are eager to boost profits through good old industrial espionage so they can copy trade secrets and underbid competitors.

In all of these cases, theft of sensitive data may tip the balance of power in an increasingly multipolar world.⁴⁷

45. <https://www.recordedfuture.com/russian-malware-analysis>


46. The Race to Build Hypersonic Missiles | U.S. vs. China | WSJ <https://www.youtube.com/watch?v=rcZwk9hmCN8>

47. <https://go.recordedfuture.com/hubfs/white-papers/digital-asymmetry-future-business-implications-balkanizing-internet.pdf>

Section 2: How to Quantify Risk

Chapter 8

Risk — to Quantify or Qualify?



We have examined the case for risk-based security and the five leading types of risk impact. In the next three chapters I would like to describe for you a process for quantifying risks. It is simpler than you might expect, and it will help you prioritize security investments and justify your decisions using language that non-technical managers understand: monetary units.

The Meaning of Risk

Presenting to a large cybersecurity audience on the topic of quantifying risk is humorous. I can see the panic in people's eyes as they read the presentation agenda slide. Merely mentioning the words "quantify" and "risk" is like shining a bright,

flashing neon sign that says “Take out your personal device now and start perusing social media.”

I’ve learned it’s always best to engage the audience early, especially when presenting on risk quantification. I start by asking the audience for a definition of risk. I never see more than one or two hands in a room of 100 people. The answers vary, but the constant theme is “damage” or “harm” — harm to a brand, damage to information systems, damage to people, and so on.

These definitions are based on the ordinary day-to-day usage of the word, but they are not nearly rigorous enough to be used as the basis of business decision-making.

Unfortunately, finding consensus on the definition of risk is very difficult. Risk is a loaded term for many in cybersecurity, and prior experiences tend to color the perception of this important concept. There are plenty of experts in risk outside of cyber. For example, enterprises in regulated industries like financial services have robust governance, risk, and compliance (GRC) and enterprise risk management (ERM) teams for calculating all kinds of risk (compliance, financial, geopolitical, and so on). However, these teams don’t apply the same analytical rigor to the operating risk from cyber threats, and rarely have cybersecurity teams tried to adopt their methods.

Where can we look for a definition of risk that will help us manage cybersecurity?

Technical bodies are not much help. In its Network and Information Security (NIS) Directive, Article 4, the European Union Agency for Cybersecurity (ENISA) defines risk as “any reasonably identifiable circumstance or event having a potential adverse effect on the security of network and information systems.” That’s an overly convoluted definition that is also partially misleading.

Defining Risk as Monetary Loss

Instead, cybersecurity organizations should adopt the definition of risk used by almost every business manager and board of directors: the potential for monetary loss.

Obviously the loss of life would be infinitely worse, but for our purposes, outside of healthcare, most industries are focused on losing money.

Risk in this context is the possibility that an event will eventually lead to reduced company profitability. A cyber event causing damage to a company's brand or reputation can be quantified. The key question is always this: how much does a cyber event ultimately cost the business?

This is a simple but powerful definition. For people who defend things, it means that every decision can be guided by the answers to three questions:

1. If we take no action, what is the risk (how much money are we likely to lose from the five risk impacts)?
2. If we take the action, how much does it reduce the risk (how much *less* money are we likely to lose)?
3. What is the cost of the action?

When cybersecurity professionals answer these questions, they speak the language of business. They can remove the cost center label and show how they are increasing profits. Their budget requests can be compared against the requests of manufacturing, engineering, marketing, sales, and every other department in the enterprise. They can communicate with executives and board members who may have little understanding of technology or security.

Cybersecurity groups can approach risk the same way insurance companies do. Those firms don't write policies without understanding the risks. They use actuarial tables to underwrite life, property, and casualty policies, and if applicants don't fall within an acceptable range on any number of variables, then the policies are denied. Enterprise executives should require the same type of analysis for their own security functions to better understand the potential for loss, and whether security control changes are required to reduce potential losses to acceptable levels.

For example, in 2009, following a cyberattack, Google began evaluating the Yubikey, a hardware authentication device in the form of a USB that uses public key cryptography to improve multi-factor authentication (MFA). The company ultimately decided to purchase and provision Yubikey devices for its workforce of more than 50,000 employees and contractors.

The technology review required two years⁴⁸ and the hardware cost for the entire enterprise was likely a big chunk of change, but Google was ahead of the curve on defeating credential reuse attacks. I don't have visibility into the conversations inside Google at the time, but the business justification was likely communicated in quantitative terms. It's unlikely that Google made the necessary investments in 2009 without quantifying not just the resource costs but also the long-term risks of inaction.

Fourteen years later, organizations are still grappling with phishing attacks, malware, and credential reuse. Actors are bypassing one-time passwords (OTP) and other MFA mechanisms. Today Google runs one of the most secure email services in the world, Gmail, and looks prescient in its long-term commitment to high levels of security.

But We *Can't* Estimate Risk in Cybersecurity (Can We?)

"That sounds wonderful," I hear you say, "but in cybersecurity it is simply not *practical* to evaluate risk in monetary terms. There is little or no historical data for the new threats we face every day. We could never construct a financial model to capture all the detail needed for those calculations. Even if we could, we don't have anywhere near the time or the staff to estimate risks and costs precisely."

I understand your concerns, but let me assure you that in cybersecurity it *is* practical to evaluate risk in monetary terms. The key tools are the systematic use of estimation, which I'll discuss here, and a practical framework for risk modeling, which I'll present in the next chapter.

The Power of Estimation

As Douglas Hubbard and Richard Seiersen point out in their seminal work "How to Measure Anything in Cybersecurity Risk," everyone wants perfect historical data for modeling, but such data is not necessary to create a meaningful model.

48. https://resources.yubico.com/53ZDUYE6/at/6r45gck4rfvbrspjxwrmcsr/Forrester_Report_Total_Economic_Impact_of_Yubico_YubiKeys.pdf?format=pdf

Hubbard and Seiersen make a compelling case for estimation; that is, training the brain to more accurately estimate values and allowing for black swan-type events. In the exercises that Hubbard and Seiersen present, the goal is a 90% confidence interval (where the correct value falls somewhere in the estimated range nine out of ten times). The estimator must be confident that the correct value falls in the range between a low and high value.

For example, unless you're a student of European history, you likely don't know the exact year that the Battle of Waterloo was fought. Without skipping ahead, think of a range that fits here. What's your best estimate? You likely know that Waterloo occurred in Europe and you may know that it involved Napoleon. When calculating a range for the Battle of Waterloo you might guess a low value of 1500 and a high value of 1900. History buffs may define a tighter range of 1700 to 1850. The Battle of Waterloo occurred in 1815. If that year falls within your range, you correctly completed the estimate exercise.

Similarly, you can build a cybersecurity risk model by estimating ranges of monetary loss due to different cybersecurity events. You don't need to know the exact loss, but rather a range of reasonable losses.

Estimation Training Is Important

Estimation exercises are important to train the brain to account for uncertainty and overconfidence. A minority of people tend to be under-confident in their knowledge; a majority of people have an issue with overconfidence when estimating ranges. Trained estimation can help fill the gaps of imperfect historical data, especially when combined with valid statistical approaches like Monte Carlo simulations, which I'll explain in a moment.

Bias in estimation is what must be acknowledged and adjusted for to create higher-quality risk model results. When my colleague Dr. Bill Ladd and I walk clients through trained estimation exercises, they are surprised and dismayed when their estimate ranges are incorrect for half or more of the first 10 trivia questions. Overconfidence causes them to supply too

narrow a range. But after multiple rounds, participants learn to widen their ranges to accommodate their lack of confidence in an answer. It's fun to watch them begin to understand their bias and adjust accordingly.

When the exercises move from random trivia to impact and loss across threat categories, the participants are rightfully wary of creating an estimate range without deep thought and consideration about their knowledge of the threat and the state of the organization's internal security controls.

Additionally, it's interesting to watch as participants factor in loss mitigation controls like cyber insurance. The deductible for a major loss event may be a million dollars, making senior executives feel comfortable capping their high-end estimate loss value at that amount, even if the insurance coverage hasn't been thoroughly tested industry-wide.

Monte Carlo Simulations

Have you heard the joke about the statistician who nearly drowned trying to cross a river? He was informed that the average depth was three feet, and was surprised to find a seven-foot drop in the middle.

A risk analysis needs to consider not only averages ("expected values" in the terminology of probability), but also unlikely but possible minimums and maximums. These include "perfect storm" scenarios, in which two or more bad things happen in the same period. A business might be able to overcome a flood, and it might be able to recover from an earthquake, but could it survive a flood and an earthquake in the same year? If not, what is the best way to reduce the maximum possible loss to an acceptable level: build a levee, earthquake-proof the headquarters building, or just buy more insurance?

Questions like those can be answered using Monte Carlo simulations. These involve selecting one random value for each model input out of a specified range and calculating the resulting losses. The simulation can be repeated thousands (or millions) of times and the distribution of losses can be examined.

Monte Carlo simulations are practical and easy to implement.

In some cases, they can be computed and updated in an Excel spreadsheet. In the next chapter, I'll explain how they can be used.

After Quantification, Communicate with Stories

I have to qualify some of my earlier statements that quantitative risk analysis is simple and provides a language that non-technical managers can understand. Those statements are true for many non-technical executives who happen to have analytical minds. But I have found that for many other executives, the findings from quantitative analyses need to be translated into binary statements and simple stories.

I believe that security leaders should still start by using the risk analysis processes described in this chapter and the next two. They enable people knowledgeable about the nuances of security to prioritize security investments and have confidence that the ones they recommend are fully justified.

However, in many (perhaps most) organizations, top executives and board members don't have the time or inclination to examine quantitative analyses. Any mention of numbers or models (beyond the simplest) leads to confusion, and potentially to questionable decisions. The typical risk matrix is often not simple enough. A heat map risk categorization may be too convoluted. Even a traffic light (red/amber/green) dashboard can cause indecision. How should a board member respond to an amber (medium) risk?

In my experience, business executives and board members are interested in simple stories. Simple means without jargon or technical details. Narratives linking together easily understood past and future events are an optimal way to communicate.

Binary statements are another effective way to communicate with executives. In a recent CISO roundtable that I attended, the consensus was that the best statements of risk are binary ones — unambiguously red or green, good or bad, better or worse, or true or false.

Unfortunately, probabilities are not binary (except for zero percent and 100 percent). However, once you have used probabilities to quantify risks and prioritize security investments, you can find ways to translate your conclusions into binary statements and, where needed, stories.

Let's look at an example of each.

First, a binary statement. "The LockBit ransomware gang has victimized two of our major competitors in the last three months. We must improve controls to reduce the risk of operational disruption and compliance failure."

The link between the threat and the recommendation is simple and compelling. While a quantification of the risk and the cost of remediation would be useful, they may not be necessary to persuade executives or board members to take the recommended actions.

Unfortunately, the world is not always that simple. Sometimes a narrative is needed to walk non-technical managers through the logic.

Here is an example of a story.

Threat actors working on behalf of North Korea (DPRK) were recently observed targeting our competitors with social engineering campaigns. Given our similarity to those competitors and the size of our reported earnings, we are a logical next target for their campaigns involving unauthorized funds transfer and financial fraud. Our risk is even higher because right now we are not very well protected from malware inserted by threat actors in products we obtain from our suppliers — exactly the kind of attacks the DPRK has been launching recently. To reduce this risk, we suggest removing the list of our vendors and suppliers from our website. In addition, we recommend establishing a threat hunting team with the specialized skills to search our network for signs of suspicious activities known to be used by agents of the DPRK.

This narrative provides a deeper explanation of cause and potential effects for board members (an example of second-order thinking⁴⁹, which I'll describe in the next chapter)

49. <https://intelligence2risk.substack.com/p/unlocking-second-order-thinking-risk-analysis>

without going too deeply into technology. Again, a detailed quantification of risk might be very helpful for some board members, but many may be more easily convinced by this type of report.

Another communication technique, advocated by my colleague Jason Steer, Recorded Future's CISO, is to highlight security wins and the progress of the security program. You can quantify how often attacks have been blocked and the savings, or the number of additional people/systems/countries covered each quarter by the rollout of a new control.

Storytelling shouldn't be an opportunity to engage in fear, uncertainty, and doubt (FUD). Rather, it involves crafting compelling narratives that highlight positive security outcomes and risks avoided from executing on a business-aligned cyber program.

As I write this in 2023, economic uncertainty abounds. Businesses are engaged in cost optimization and vendor consolidation in the midst of rapid and consequential global changes — conflicts, wars, recessions, “slowcessions,” generative AI, digitalization, hybrid work, and more. As organizations strive to optimize operations and correct past excesses, security leaders have an opportunity to tell a story of optimization while balancing risk management perceptions. Stories, backed up by facts and solid analysis, create confidence — and confidence is what executives and boards of directors need.

I am also encouraged by the fact that the U.S. Securities and Exchange Commission and other regulatory bodies are proposing that boards of directors include at least one member with cyber experience. That may result in more boards with members who can help raise the security literacy of their peers.

A Last Resort

After working with Recorded Future clients for the past five years, I have learned that it can be very challenging to win the hearts and minds of traditional risk departments. When ERM/GRC groups or executives resist the concept of cyber risk quantification, even the most compelling quantitative analyses and presentations will fall flat. And although quan-

titative conversations are more precise and instructive, the majority of leaders I encounter are most comfortable with simple stories.

But I have also learned that even if an enterprise refuses to move forward with cyber risk quantification, you can achieve a lot by deploying binary statements, qualified risks, and stories backed by solid intelligence. At the end of the day, security professionals who have to operate with pre-existing communication styles can still use risk quantification and qualification to better protect their business.

Chapter 9

Quantifying Risk with the Threat Category Risk Framework



We mentioned earlier that one requirement for risk-driven cybersecurity is a practical framework for risk modeling. Unfortunately, the best-known cyber threat taxonomies and frameworks, the Diamond Model,⁵⁰ the MITRE ATT&CK Matrix,⁵¹ and the Lockheed Martin Kill Chain,⁵² although helpful tools, are oriented toward identifying and remediating threats, not risks.

As with compliance frameworks, the population of cyber threat models should never represent the end-state goal of a security team. If a tool or framework is too convoluted and

50. <http://www.activeresponse.org/wp-content/uploads/2013/07/diamond.pdf>

51. <https://attack.mitre.org/matrices/enterprise/>

52. https://www.lockheedmartin.com/content/dam/lockheed-martin/rms/documents/cyber/Gaining_the_Advantage_Cyber_Kill_Chain.pdf

not practical to use, then consider building a framework that is better suited for the human resources available and the desired outcomes. Framework categories should be intuitive and segmented at a reasonable level of granularity. “Practical” is obviously a subjective characterization, but like Supreme Court Justice Stewart’s test of what constitutes pornography, cybersecurity professionals should know it when they see it.

For example, I’ve observed security teams that spend months mapping one threat group’s tactics, techniques, and procedures to the ATT&CK framework. That exercise helped improve internal network hunting methodologies. However, the time spent mapping that one group created a deficit of understanding for the tactics, techniques, and procedures (TTPs) of hundreds of other threat actors. In other words, the time spent on overly granular mapping isn’t worth the benefit, especially when human resources are limited.

Similarly, deliberating over whether an adversary technique falls into the Kill Chain’s “Phase 3 — Delivery” or “Phase 4 — Exploitation” is counterproductive. What’s important is surfacing techniques and assessing them against existing security controls.

The Diamond Model focuses on mapping adversary infrastructure and capabilities as they relate to a victim. This model is especially helpful when attempting to attribute malicious activities to adversaries. However, it’s less helpful outside of the public sector, where attribution is the operational outcome.

Introducing the Threat Category Risk Framework

The threat category risk (TCR) framework, built on Hubbard and Seiersen’s work, is a practical, quantitative risk framework designed to clearly articulate the probability and amount of economic loss that an organization faces from cyber threats in a given year. This makes it an ideal framework to drive a risk-based security program.

The approach is very simple. The TCR framework starts with a set of general threat categories. For each threat category, a team estimates:

- The “event risk,” which is the probability the event will occur in the coming 12 months
- The probability that, if the event does occur, it will result in the loss of confidentiality or integrity (that is, the improper disclosure of information or the unauthorized modification of data or system behavior), or the loss of availability (that is, a system outage), or both.
- The upper and lower bound of damage if a loss of confidentiality or integrity occurs
- The upper and lower bound of the duration (in hours), and the upper and lower bound of the cost per hour if a loss of availability occurs

Based on these estimates, a relatively simple calculation will reveal not only the most likely loss, but also a range of possible losses from the threat category.

We will walk through an example of the calculation in a moment, but you can probably grasp already a couple of significant characteristics of the TCR framework:

1. It calculates risk in monetary terms.
2. A team with the right skills, knowledge, and training in estimation should be able to provide the inputs with a reasonable amount of accuracy (especially because several of them are ranges) in a reasonable amount of time.

The TCR Threat Categories

The first step in using the TCR framework is to select the threat categories that are relevant to your enterprise.

The TCR categories listed in the table below are general on purpose. For ease of use and simplicity, they are divided between initial compromise methods and post-compromise methods — sometimes called “left of boom” and “right of boom,” respectively. As Benji Hutchinson explained: “Popularized in military circles during the months and years after 9/11, the phrase ‘left of boom’ refers to the moments before an explosion or attack — a period when you still have time to prepare and avert a crisis. Right of boom, by contrast,

includes the chaotic and deadly moments after the explosion or attack.”⁵³

Note that TCR avoids excessive granularity in attack types, because great precision in estimating impact and loss ranges is not necessary in this framework. That saves us a lot of time and effort, because we only need to estimate the probabilities and impacts for a few threat categories.

Initial Compromise (Left of Boom)	Post-Compromise (Right of Boom)
Social Engineering*	Denial of Service (DoS)
Credential or Key Reuse/ Stuffing/Brute Forcing	Theft of Employee or Customer Personally Identifiable Information (PII)
Misusing Open Ports/ Network Shares (Manual or Automated — Worms)	Theft of Proprietary Communications or Information
Web Application Vulnerabilities (Including Web Shells)	Access and Theft of Data from Connected Third Parties
Hardware Vulnerabilities	Blackmail/Extortion
Software Vulnerabilities	Destruction of Data or Systems Availability
Protocol Hijacking (BGP/DNS)	Removal of Confidence in Data Integrity
Physical Tampering	Financial Fraud

*Includes phishing, spear phishing, business email compromise, and mislabeling malicious files in P2P networks

The primary difference between TCR and other frameworks is that the threat categories are aligned to monetary loss. TCR isn’t an adversary-centric framework, like the Diamond Model, because that would be redundant — it’s implied that an adversary is manually or programmatically launching attacks.

Also, we don’t have to analyze every possible threat category. We can focus on those that directly cost the business money. Some adversary tactics are important to detect because they indirectly contribute to loss, but for the purpose of calculating potential economic losses, they are less relevant. TCR is concerned with the threat categories and the subsequent actions that cause loss of confidentiality, integrity, and availability of systems and data.

53. <https://necteday.com/left-of-boom-defeating-the-threat-among-us/>

Walking Through an Example: Credential Reuse

Let’s walk through the process of making estimates for one threat category that affects Acme Corporation: credential or key reuse/stuffing/brute forcing (we’ll call it “credential reuse” for short).

Credential reuse typically occurs when an attacker steals credentials during a data breach (or purchases them on the dark web) and tests them against many websites and social media accounts. It’s a very effective and inexpensive way to penetrate networks and gain access to confidential data and IT resources.

So how would we go about estimating the risk of credential reuse? The following charts are based on Hubbard and Seiersen’s work.

We start by estimating the likelihood that the event will occur within the next 12 months. For simplicity, the second column in the table below (“Event Risk”) is summarized as a percentage instead of a high/low range estimate, but when implementing this model it’s worthwhile to create range estimates for event risk as well.

Risk Type	Event Risk	CI Only	AV Only	Both
Credential Reuse	100%	60%	30%	10%

If a threat category is relevant (the event risk is above zero), the next step is estimating, if the event occurs, how often it will affect confidentiality and integrity only (“CI Only”), the availability of data only (“AV Only”), or both. For example, if I’m estimating values for Acme Corporation, I might estimate that the credential reuse threat category will impact information confidence/integrity only 60% of the time, availability only 30% of the time, and both 10% of the time.

If data confidentiality/integrity are impacted by a threat category (row), then the CI Low and CI High columns must be

populated with a low-value estimate and a high-value estimate of cumulative losses in the next 12 months. For the Acme example, I estimate that incidents involving credential reuse will cost no less than \$1,000 and no more than \$25,000 over the next 12 months.

Risk Type	Event Risk	CI Only	AV Only	Both	CI Low	CI High
Credential Reuse	100%	60%	30%	10%	\$1,000	\$25,000

I’ve also determined, for the purposes of this exercise, that credential reuse is a threat category that could affect both data confidentiality/integrity and data availability. When data availability may be affected, I must provide low- and high-value time estimates (columns “Time Low” and “Time High”) that are again aggregated for an annual time period.

Risk Type	Event Risk	CI Only	AV Only	Both	CI Low	CI High	Time Low	Time High	AV Low CPH*	AV High CPH*
Credential Reuse	100%	60%	30%	10%	\$1,000	\$25,000	50.0	300.0	\$100	\$250

*Cost per Hour

Related to the credential reuse threat category, I estimate that my low boundary for the year is 50 hours of lost data availability, and my high boundary is 300 hours. To create these estimates I need to understand the basic capabilities of current credential reuse TTPs used by attackers and any mitigation controls that are in place to defend against them. Lost data availability may come from attackers’ activities like misconfiguring network and security devices, shutting down servers, or destroying hard drives via wiper malware.

Finally, the “AV Low CPH” and “AV High CPH” columns represent low and high dollar amount estimates for the cost per hour due to the unavailability of systems or information. For example, if critical applications are not available due to an attack on a key server, I might estimate that the company will lose between \$100 and \$250 per hour.

Putting the Categories Together

One advantage of the TCR framework is that it can be created and maintained on a spreadsheet (possibly a cloud-based spreadsheet for easy sharing and editing).

Here is how our matrix might look if I entered estimates for all of the threat types facing ACME Corporation.

Risk Type	Event Risk	CI Only	AV Only	Both	CI Low	CI High	Time Low	Time High	AV Low CPH	AV High CPH
Social Engineering	100%	80%	10%	10%	\$ 20,000	\$ 250,000	1,200.00	3,000.00	\$ 100	\$ 250
Credential Reuse/Stuffing/Brute Forcing	100%	60%	30%	10%	\$ 1,000	\$ 25,000	50.00	300.00	\$ 100	\$ 250
Web Application Vulnerabilities	50%	40%	30%	30%	\$ 5,000	\$ 500,000	24.00	120.00	\$ 500	\$ 100,000
Denial of Service	5%	0%	100%	0%	-	-	1.00	24.00	\$ 70,000	\$ 240,000
Internet Protocol Hijacking (DNS/BGP)	5%	10%	80%	10%	\$ 100	\$ 1,000,000	.50	72.00	\$ 10,000	\$ 250,000
Hardware Vulnerabilities	50%	10%	60%	30%	\$ 1,000	\$ 1,000,000	3.00	336.00	\$ 500	\$ 100,000
Software Vulnerabilities (not web related)	100%	0%	100%	0%	-	-	50.00	1,500.00	\$ 100	\$ 500
Physical Tampering	10%	0%	90%	10%	\$ 100	\$ 1,000,000	.25	168.00	\$ 500	\$ 50,000

Running the Monte Carlo Simulation

Now that I've input my estimates for each threat category, I can run a Monte Carlo simulation and output the resulting median values for each row (shown below).

I might specify 100,000 simulations, but since increasing the simulation count doesn't alter the median value variation significantly, there's no reason you can't specify one million simulations if you're so inclined.

The simulation will give us insight into both:

- The expected value of the loss for each risk category
- Unlikely, but potentially catastrophic, outcomes we might want to guard against

The first column of this spreadsheet shows the probability that a certain loss or a greater one will occur. For example, if you look in the "Total Loss" column, you can see that there is a 50% probability that Acme will lose \$2.3 million or more, and a 1% probability that it loses \$31 million. Likewise, if you look in the "Credential Reuse: Total Loss" column, you will see that in 40% of the simulations Acme loses roughly \$1.1M or more in the next year from that threat category.


Most organizations understand basic losses represented in the top half of the percentile chart. Looking at the 50% row and

seeing that the “expected” total loss is around \$2.3 million, management might feel quite comfortable with the status quo (especially since the cost of operating an incident response team to triage successful phishing incidents or commodity malware infections can easily cost north of \$2 million a year in employee compensation and technical tools).

Probability of This Loss or Greater	Total CI Loss	Total AV Loss	Total loss	Credential Reuse: Total loss	Web Application Exploitation: Total loss	Exploited Vulnerability: Total loss	Phishing: Total loss	Ransomware Internal Workstations Only: Total loss
95%	\$48,932.98	\$1,108,327.00	\$1,251,663.00	\$580,034.60	\$0.00	\$0.00	\$217,901.40	\$9,404.66
90%	\$63,457.61	\$1,242,980.00	\$1,403,847.00	\$657,155.40	\$0.00	\$0.00	\$247,986.70	\$14,276.62
85%	\$75,957.43	\$1,350,330.00	\$1,522,604.00	\$714,480.80	\$0.00	\$0.00	\$270,497.70	\$18,792.91
80%	\$87,625.05	\$1,444,638.00	\$1,630,886.00	\$763,884.80	\$0.00	\$0.00	\$289,878.00	\$23,510.41
75%	\$99,539.46	\$1,534,862.00	\$1,734,583.00	\$809,716.70	\$0.00	\$0.00	\$307,783.70	\$28,501.25
70%	\$111,585.25	\$1,625,478.00	\$1,838,824.00	\$852,792.00	\$0.00	\$0.00	\$324,795.00	\$33,866.94
65%	\$124,293.34	\$1,719,219.00	\$1,945,519.00	\$894,525.00	\$0.00	\$0.00	\$341,176.70	\$39,657.47
60%	\$137,632.16	\$1,817,014.00	\$2,059,930.00	\$935,650.50	\$0.00	\$0.00	\$357,726.00	\$45,961.65
55%	\$152,428.25	\$1,924,844.00	\$2,182,548.00	\$977,639.10	\$0.00	\$0.00	\$374,549.50	\$53,213.17
50%	\$168,940.79	\$2,044,278.00	\$2,320,317.00	\$1,020,942.60	\$14,674.47	\$0.00	\$391,791.10	\$61,335.51
45%	\$187,401.55	\$2,181,988.00	\$2,477,782.00	\$1,064,832.10	\$101,975.60	\$45,128.89	\$409,867.10	\$70,839.06
40%	\$209,070.60	\$2,341,963.00	\$2,666,050.00	\$1,113,015.90	\$174,960.29	\$93,787.96	\$449,749.10	\$82,037.47
35%	\$234,614.25	\$2,538,985.00	\$2,892,137.00	\$1,165,631.10	\$260,710.40	\$160,561.43	\$494,759.60	\$95,359.09
30%	\$266,081.91	\$2,794,515.00	\$3,178,194.00	\$1,223,398.50	\$373,392.09	\$254,378.97	\$473,006.80	\$111,632.72
25%	\$306,010.38	\$3,136,882.00	\$3,570,372.00	\$1,288,149.30	\$524,224.21	\$395,223.73	\$499,291.50	\$132,252.68
20%	\$361,115.69	\$3,633,943.00	\$4,110,297.00	\$1,363,878.70	\$747,733.01	\$616,575.30	\$530,187.90	\$160,355.11
15%	\$442,934.33	\$4,425,370.00	\$4,966,129.00	\$1,459,628.40	\$1,094,014.58	\$1,010,318.04	\$569,239.00	\$200,229.88
10%	\$586,292.40	\$5,885,030.00	\$6,496,379.00	\$1,589,005.90	\$1,754,853.29	\$1,828,368.38	\$622,563.20	\$265,061.73
5%	\$948,516.43	\$9,657,421.00	\$10,436,777.00	\$1,801,617.20	\$3,451,042.16	\$4,297,746.59	\$710,808.00	\$401,099.91
1%	\$2,984,718.46	\$29,677,965.00	\$31,066,102.00	\$2,289,354.70	\$11,883,288.27	\$20,592,835.66	\$918,448.80	\$871,851.70

However, we also need to consider the figures below the fiftieth percentile row, particularly those that document the probability of loss between 50% and 15% (shaded in the figure). These loss amounts should encourage conversation. An organization that would accept a 50% probability of a \$2.3 million annual loss may reject a 15% probability of losing \$5 million in the current year.

Ultimately, TCR should generate a loss exceedance curve like the one below that can be featured prominently in communication with the board of directors (assuming the board members have appetite for risk quantification).



Example of a loss exceedance curve (Source: Paul Stokes articles on the World Economic Forum website: <https://www.weforum.org/agenda/2019/07/can-cybersecurity-offer-value-for-money/>.)

The advantages of the TCR framework include simplicity, transparency, minimal resource requirements (a spreadsheet), and practicality (one or two days to train estimators on the process of estimating ranges for their organization). Because the model inputs of estimated ranges of loss are clearly specified, they can be discussed and improved if better estimates become available.

Estimating the Value of New Security Controls

In addition to estimating the probability of loss, the TCR framework enables security practitioners to estimate the value of implementing new security controls. Based on the expected improvements in security from the new controls, they can change inputs for the probability of event occurrences and the range of losses. They can then generate a new set of probable losses, calculate the delta, and compare the projected savings with the cost of the controls. The result is a dollar figure that can be shared with executives and accountants alike, using their own language to justify the investment.

Complementary to Control Validation

One additional method for consistently communicating security control improvements is control validation platforms. These are iterative approaches to testing security controls against realistic attacks such as red team exercises. Control validation is a beneficial tool for communicating security changes. It can complement quantitative risk scoring and sometimes even replace it in cases when security teams aren't ready to fully embrace TCR or other quantitative risk frameworks.

Chapter 10

Updating the Framework: RTDs and Threat Intelligence



Relevant Threat Deltas (RTDs)

In the previous chapter we introduced the threat category risk (TCR) framework and discussed how to create an initial model for your organization. But once you have your model, when and how do you update it?

You might think the estimates need to be updated frequently. After all, cyberattacks occur every week, if not every day (or even every minute — have a look at a live perimeter firewall log). However, basic security controls will obviate the impact of attacks on most businesses, and truly innovative new threats are rare. Adversary technical innovation is largely opportunistic, and the pool of adversaries with advanced skills

and an ability to innovate is relatively small when compared to the total pool of actors who are active in the underground economy.

However, you do need to update your TCR model when what I call a “relevant threat delta” (RTD) occurs. RTDs are caused by:

- New or modified threats sufficiently innovative to evade existing controls
- Changes to the organization’s business or technology that expose it to threats that were not previously relevant

In other words, RTDs are events that change risks enough to have a material impact on your TCR model.

In practice, RTDs are infrequent enough so that TCR frameworks only need to be updated quarterly, bi-annually, or even annually.

That being said, you must be vigilant to ensure that you do discover RTDs in a timely manner. Businesses that don’t make quick security control adjustments in response to changing threats are at increased risk of monetary loss. When a new RTD is discovered, there is a gap between the new threat and a business’s security control response (which may include an information security vendor’s updated response). Eventually the business or related security vendor will catch up, but businesses must be wary about those windows of adversarial opportunity.

Why Threat Intelligence Is Critical for Risk-Based Cybersecurity

If RTDs initiate changes to your TCR framework, which drives your cybersecurity program, then several questions follow:

- How do you discover RTDs in a timely manner?
- How do you separate *relevant* threat deltas from *minor* threat deltas and *relevant-for-others-but-don’t-affect-my-organization* threat deltas?

- How do you know how much to adjust your risk estimates when you verify that an RTD has occurred?
- Perhaps most important, how do you find the best options for improving security controls to minimize the impact on your security posture?


The answer to these questions is threat intelligence, especially strategic threat intelligence.

Threat intelligence gives you visibility into the threat environment, including information about active adversaries and their TTPs. It enables you to discover RTDs early, sometimes in the planning or development stage, before attack campaigns are launched.

Also, by comparing adversary TTPs against your existing security controls, you can make informed judgments about what threats are relevant and material to your organization, the potential effects on your risk profile, and possible countermeasures.

The diagram below illustrates the process for keeping your TCR framework up to date:

1. Threat intelligence allows you to identify threat events that change your inputs to the TCR risk model (the RTDs).
2. Your estimators use descriptions of the RTDs and related threat intelligence to update their probabilities and estimates.
3. You rerun the Monte Carlo simulations with the revised inputs to produce new monetary estimates of risks.
4. The outputs of the Monte Carlo simulations allow business managers and cybersecurity professionals to work together to make decisions about acceptable risks and changes to security controls.



Section 3: Intelligence and Risk Management for Business

Chapter 11

Strategic Threat Intelligence



The Value of Intelligence

As we mentioned in the previous chapter, threat intelligence enables businesses to identify adversary tactics, techniques, and procedures (TTPs) and determine whether new TTP instances will render existing security controls insufficient. Threat intelligence, via RTDs, should drive risk score changes, or measurably improve operational security, or do both.

This chapter addresses the ingredients of a good threat intelligence program, and the direct benefit to the business in tangible terms that demonstrate decreasing operational risk of economic loss through better security.

What Is Intelligence?

Intelligence constitutes a significant and growing chunk of business security budgets. A recent survey found that 85% of IT organizations are currently using a threat intelligence service or planned to start using one within 12 months.⁵⁴ But what exactly *is* intelligence?

My definition of intelligence is the act of formulating an analysis based on the identification, collection, and enrichment of relevant information.

Analysis is the key. It is the bridge between information and intelligence. Analysis is only accomplished through the separate and combined effort of the left and right sides of a human brain (or well-trained machines). The process and result of intelligence comes in many forms and applications.

In the professional world, intelligence is applied to a myriad of business problems, one of which is adversaries that seek to disrupt the confidentiality, integrity, and availability of information belonging to their victims. This is also known as a “threat.” A practical definition of threat intelligence is defensive improvements created through analysis of the adversary’s operating space.

Intelligence provides an information advantage to connected enterprises. Since the beginning of time, humans have been seeking an edge. That pursuit has evolved through history. Today we all seek an information advantage in our daily lives — in sports, in traffic on our way to work, when shopping for a new car or buying groceries. How momentarily excited are you when your phone suggests an alternate route to work that saves you 10 minutes or you discover a website selling the same product for \$50 less?

Intelligence Leads to a Persistent Decision Advantage

Threat intelligence allows organizations to anticipate risks and either head them off or react before they cause significant loss. It also gives enterprises and law enforcement a shot at attrib-

54. <https://go.recordedfuture.com/cyberedge-cyberthreat-defense-report-2019>

uting attacks to their perpetrators, which changes the odds of catching the bad guys.

Early in my career with the U.S. Secret Service, I remember picking up my desk phone and taking a report from a man who explained how his mother had been victimized by a Nigerian email scam to the tune of half a million dollars. Naturally, this man was upset, and I felt terrible for him and his mother. Sadly, the likelihood of recovering funds at that time was slim to none, and slim was walking out the door. I quickly realized that threat intelligence was necessary to develop quality criminal leads that proactively generated cases before a victim picked up the phone.

Without intelligence and significant resources to pursue attribution, it's difficult to solve a cybercrime case. The biggest cases take years to prosecute. For the good guys, it feels like a game of whack-a-mole, with new criminals springing up quicker than the old ones can be nailed.

Start by Strengthening Basic Security Controls

Basic security controls are a good litmus test for more-advanced security measures like threat intelligence programs, just as door locks and a knowledge of where doors and windows are located are prerequisites for a home security system.

For instance, before the security operations center (SOC) and incident response (IR) functions can work effectively, it is necessary to generate, collect, and analyze comprehensive host and network-based logs. Collecting breached credentials from criminal forums and automating the process of password resets are examples of valuable new security controls, but from a basic risk perspective, priority should be given to addressing shortcomings in password complexity and storage requirements.

Don't Rely on Daily Threat Reports

I have found that threat intelligence leaders often make the mistake of hiring analysts to create daily threat reports to increase awareness of threats throughout the business. In my

experience, that's rarely a goal worthy of the budget necessary to create the capability.

A few years back, I was on site with a Fortune 500 client and asked about their threat intelligence program goals and the associated deliverables. The answer to both questions was a daily threat report. I asked how the reports created operational outcomes and how those outcomes were measured and communicated. I received a room full of shrugs.

That's about the time that I registered a small explosion in my brain. These were talented and motivated analysts working for a premier global enterprise, but their role had been reduced to secondhand reporting for the purpose of increasing awareness. I wanted to rewrite their mandate and charter on the spot, but of course that was beyond my control.

Daily Threat Reports Versus Useful Reports

Here's the difference between a topical daily threat report and a less-periodic, more-extensive report that includes assessment details.

A daily threat report is typically a short, bulleted list of facts obtained from threat databases, together with associated impact ratings. A section might look like this:

- BreachForums shuts down following the arrest of “pompompurin” (impact: low)
- BlueBravo uses ambassador lure to deploy GraphicalNeutrino malware (impact: low)
- Valiant Panda: China's use of political cartoons in malign influence campaigns targeting US (impact: medium)
- Supply chain attack on business phone provider 3CX could impact thousands of companies (impact: high)
- IceFire ransomware Linux variant targets media and entertainment sector (impact: medium)
- Samsung had its internal and confidential data leaked after using ChatGPT (impact: high)

Compared with a daily threat report, a useful threat report contains more valuable, in-depth analysis that drives operational outcomes:

- On March 15, 2023, a critical elevation-of-privilege vulnerability (CVE-2023-23397) affecting all supported versions of Microsoft Outlook was published. The attacker can send a malicious email to a vulnerable version of Outlook to obtain the Net-NTLMv2 hash, which can be used to authenticate to other services by impersonating a victim without user interaction. On March 17, the vulnerability team patched all vulnerable servers.
- On May 4, 2023, a new malware variant, FluHorse, was discovered. It disguises itself as well-known and legitimate Android applications to steal victims' sensitive information such as credentials, credit card data, and two-factor authentication (2FA) codes. The mimicked applications are the Taiwanese Electronic Toll Collection application (ETC) and the Vietnamese bank application (VPBank Neo). We've sent out communications to all employees (including those in Vietnam) to remind them of our existing security policy, which states that all new applications should come from certified app stores (like Apple's and Google's stores).
- On February 13, 2023, we identified a newly registered domain with lexical similarity to one of our brands. On February 17, 2023 we observed the creation of a new DNS record resolving to the typosquat domain. Additionally, on February 17, 2023, the typosquat domain's associated web server began using a self-signed SSL certificate. On March 1, 2023, we issued a domain takedown request with a third-party service.
- Between April 2 and April 10, 2023, we identified five new compromised employee credentials. We generated Active Directory password resets for the affected employee accounts.
- On May 10, 2023, researchers warned defenders about a new phishing-as-a-service (PaaS) tool allowing rookie hackers to incorporate "some of the most advanced" features into their cyberattacks. A few of

these features include multi-factor authentication (MFA) bypass, IP filtering, and integration with Telegram bots. Our corporate security team has enforced the use of hard-key MFA (Yubikeys) for all employees with access to critical data and assets.

The difference between the content in these two example reports is stark. The first report contains references to events that have been previously reported elsewhere. The impact estimates it provides are based on guesswork; its categorizations might give readers a general sense of the risk, but offer little rigor behind designations like “high/medium/low” or color terms like “red” and “amber.” As a reader, should you lose more sleep over a “medium” or an “amber” designation?

Conversely, the second report contains first-party, original reporting on new events. Those events are accompanied by thorough assessments within the context of existing security controls. Further, each bullet reports remediation status and actions leading toward a final disposition that would ameliorate risk. The second report specifies outcomes that are measurable and communicates them in language a business manager would understand.

To create the second type of report, you need both talented people and the proper tools. Even with adequate resources, certain assessment workflows require time. Given the frequency of daily threat events, it’s next to impossible to provide a valuable daily threat report without enormous resources.

For these reasons, I recommend in no uncertain terms that executives should remove the daily report requirement and direct analysts to focus on quality regardless of cadence. Quality always trumps volume in private sector intelligence reporting.

Don’t Create Reports for Nonexistent Audiences

Beyond occasional relevance for the public relations or legal department, threat reports that lack detailed security control assessments are in danger of serving a nonexistent audience. If you’re building or managing a threat intelligence capability

in which the primary deliverable is reporting, you should ask yourself:

- Who reads these reports?
- How do the reports impact business decisions, particularly around security spending?
- What operational outcomes are occurring as a result of these reports?
- How do we measure and communicate the outcomes produced by reporting?

Stop Saying “Actionable”

There’s a lot of confusion about what “actionable” means, although it’s a popular word to throw around in meetings with executives. When I talk about threat intelligence with partners and clients, they often say, “I need intelligence that’s actionable.” That leaves it up to others in the room to interpret their intent, which usually produces unexpected outcomes.

To be actionable, intelligence must have certain criteria that can be measured in consistent, unambiguous units understandable to the intended audience. That kind of intelligence can:

- Cause changes in our systems, processes, or workflows
- Be measured in concrete ways, for example by changes in threat modeling, security controls, productivity, or costs
- Be communicated in a language that the audience understands, whether it is the rest of the security team, a manager, or an organization’s board of directors

Risk-Based Analysis Helps Threat Intelligence Produce Operational Outcomes

I said earlier that threat intelligence is essential to a risk-based security program because it enables you to discover RTDs and make informed judgments about what threats are relevant for

your organization, the potential effects on your risk profile, and possible countermeasures.

But the reverse is also true: risk-based analysis like that provided by the TCR framework is needed to use intelligence effectively.

In many organizations, threat intelligence reports are read by a few SOC analysts and executives. However, typically nothing happens because there is no risk-based analysis that justifies the required resources, and because the threat intelligence team never finds out why its output is ignored.

Let's say the threat intelligence team writes a report on the security merits of upgrading thousands of workstations from Windows 10 to Windows 11. The CIO and several direct reports read the report and decide that current cyber threats pose a risk to the Windows 10 status quo. The CIO recommends upgrading to Windows 11 and cites the threats listed by the threat intelligence team. However, the CFO makes a business decision to defer the upgrade with its million-dollar cost. In this scenario, the threat intelligence was valid but produced no operational outcomes.

If the threat intelligence team had used the TCR framework, it might have produced an analysis indicating that upgrading to Windows 11 would reduce risk by more than a million dollars in the first year. The analysis might even have shown that introducing additional security controls, or moving desktop processing to a cloud environment, would produce even greater reductions in risk and financial savings. In this case, the risk-based analysis would have pointed to positive operational outcomes and made the threat intelligence actionable for the organization.

For a second scenario, suppose a gaming company is thinking about moving operations from Las Vegas to Macau. The threat intelligence team writes a geopolitical risk report. The CEO reads the report but doesn't understand the implications of the findings, and so ignores key recommendations.

Should a business implement a cloud-access security broker (CASB) solution?⁵⁵ Should it invest in software to prevent

55. <https://www.csoonline.com/article/3104981/cloud-security/what-is-a-cloud-access-security-broker-and-why-do-i-need-one.html>

executives from being victimized by a business email compromise (BEC)? Great questions, but it is difficult to prove the value of these solutions without monetary analysis. Increased awareness by itself is a goal with no outcomes, and by extension, it offers no value that can be measured and communicated.

Sourcing

A valuable threat intelligence program sets a goal of discovering RTDs and communicating them to the proper stakeholders. From that start, it works backward, via the intelligence life cycle,⁵⁶ to create the capability for delivering the kinds of reports that actually drive decisions. Part of this process involves specifying data collection and sourcing requirements, as well as applying the human skill sets necessary to maximize the data's value.

For each of the TCR threat categories in your model, you need to evaluate data requirements. There are six broad types of threat intelligence data:

1. Open source
2. Closed source
3. Passive telemetry
4. Active telemetry
5. Customer telemetry
6. Malware-processed metadata

It's important to understand each data type and how it's collected. It may be easier, save time, and limit legal risk to use vendors or other third parties for collecting specific types of data, and each threat category may require different datasets to fulfill the collection requirements.

Open Source

The largest collection of open source data typically originates on the World Wide Web, but sources also include chat forums like Internet Relay Chat (IRC) networks, WhatsApp, and

56. <https://www.recordedfuture.com/threat-intelligence-lifecycle-phases>

Telegram. If the data is discoverable and free to collect, then it's open source data. For example, although Tor sites (sites using a .onion TLD) are often lumped under the "dark web" label and assumed to count as closed sources, unless a Tor forum requires vetting or payment to participate, the data collected there is open source.

Closed Source

Closed source data requires special access. The underlying data inhabits the same media as open source data — web, chat, and so on — but access must first be established. In the case of criminal forums, often a payment is required or members must vouch for the online moniker before access is granted. The marketing departments in a lot of cybersecurity organizations like to refer to this data as originating from the "dark web," but if vetting is required, then "closed source" is a more accurate description.

Passive Telemetry

The best way to think about passive data collection (telemetry) is to visualize a sensor or network of sensors that log interactions with other devices. A good example is a honeypot (a computer that deceptively mimics services) or "dark" IP space (darknet) that has no legitimate purpose beyond interacting with or logging activity from internet (or internal network) hosts. These collect and log packets and files from rogue hosts (and only from rogue hosts, because legitimate hosts wouldn't be interacting with dark IP space or a honeypot). GreyNoise is an example of a commercial service for passive telemetry.

Active Telemetry

Active telemetry involves scanning internet hosts and enumerating their ports, associated services, vulnerabilities, and so on. Shodan, Censys, and Binary Edge are classic examples of commercial services that actively crawl the internet and store the resulting data for customer querying.

Customer Telemetry

Customer telemetry is the data produced by a customer's endpoints or network. That data is sent to an appliance

or software owner. For example, Microsoft produces the world's most prolific software operating system, Windows. Hypothetically, if Windows collects basic system information (such as geographic locations where Windows is installed) and sends that data back to Microsoft, then Microsoft is generating customer telemetry, in part to help it improve its products. Customer telemetry is a rich source of information from large enterprises because of their access to and insights about global endpoints and networks.

Malware-Processed Metadata

Malware-processed metadata is its own threat intelligence data type because the number of malicious code (malware) samples is large. The exact number is impossible to pinpoint at any given time, but the volume is immense — somewhere on the order of yottabytes. Open source tools (such as Cuckoo Sandbox) and commercial ones (Joe Sandbox) detonate malicious code — that is, they execute files on an isolated computer or phone or in an emulated environment — and extract metadata about the actions of the malware file. The commercial services that store and analyze patterns in malware metadata are useful resources for establishing ground truth about a particular file. The best-known commercial services for malware metadata storage and searching are VirusTotal and ReversingLabs.

Private Sector Sources

Government intelligence and defense agencies have enormous intelligence-gathering capabilities, and share much valuable information about cyber threats with both public sector and commercial organizations. But the massive, non-stop avalanche of digital data exceeds the coverage of any one organization. As an illustration, China has 100,000 analysts working on interpreting open source data and producing intelligence from it.⁵⁷

Today, everyone needs to share the burden of collecting and analyzing threat information, and many private sector organizations are stepping up and playing a major role.

57. <https://www.wsj.com/articles/rise-of-open-source-intelligence-tests-u-s-spies-11670710806>

These sources of intelligence include:

- Threat intelligence providers (such as Recorded Future)
- Vendors of cybersecurity solutions that collect telemetry from their products in the field
- Cybersecurity consultants and managed security service providers (MSSPs) that monitor networks and respond to threats
- Security groups at global organizations, particularly in industries like finance, e-commerce, technology, and communications, which battle threat actors every day

Intelligence produced by private sector entities is now essential for every type of organization, even defense agencies.⁵⁸ Security teams should make a systematic effort to access the broadest range of sources possible, either directly or through intelligence aggregators like threat intelligence providers and MSSPs.

Staffing and Community Support

Strategic threat intelligence programs thrive when they are staffed by analysts with diverse skill sets.

Broadly speaking, I see analysts with three types of experience as contributing the most to threat intelligence programs:

- Military and intelligence backgrounds
- Law enforcement experience
- Technical backgrounds

Analysts from military and intelligence agencies understand the process of data collection, analysis, and reporting. They understand biases and seek clarity in their conclusions. There are private sector threat intelligence teams that dedicate whole teams of analysts to each of the intelligence life-cycle functions.

58. <https://www.foreignaffairs.com/world/open-secrets-ukraine-intelligence-revolution-amy-zegart>

I've personally observed massive teams within financial services companies that likely rival the intelligence capabilities of small countries. They have large teams of analysts and engineers dedicated to intelligence collection, analysis, and reporting.

Law enforcement analysts and agents may be less familiar with the traditional intelligence life cycle, but they have knowledge and experience about criminal tactics and methods and are accustomed to distinguishing fact from opinion.

Technical information security practitioners are critical to the successful production of threat intelligence because they have a deep background in technical security disciplines such as security operations, incident response, security engineering and architecture, vulnerability management, and red teaming. Practitioners with a technical background are necessary for their deep knowledge of security controls and offensive tradecraft, and also because they understand specialties like malware reverse engineering, infrastructure design and maintenance, and network and host-based forensics.

Only a team with multiple types of human resources can produce high-quality strategic threat intelligence. Identifying RTDs requires intelligence analysts and technical engineers to work together to discover new cyber threats, assess their impact on existing security controls, and estimate the resulting change in risk.

It's important for CISOs to support their threat intelligence team's participation in conferences, events, email lists, Slack channels, IRC channels, and other spaces where security professionals network with each other and discuss common challenges and solutions. Because cyber threats evolve quickly, it's critical for threat intelligence professionals to have buy-in from team leaders to spend time and budget on participation in communities that will benefit the security group and ultimately the business.

These communities are vital to creative and effective solutions. A strategic threat intelligence practice with continuous input and feedback from peers in similar and dissimilar industries will have a more informed and more effective team.

Avoid Siloing

You need to consider the placement of the threat intelligence function within your larger security organization. Because threat intelligence is often the new kid on the block, long-term success is dictated by the reception it receives from other security teams. On multiple occasions, I've witnessed dysfunctional enterprise security teams whose members actually view threat intelligence as a threat due to perceived role and responsibility overlap. I've personally had team leaders in lateral security groups tell me (once during my first week on the job) that they have zero interest in collaboration because they don't want to see their mission or span of control eroded.

The easiest route to continued security control improvement is to embed the threat intelligence team in a veteran group, such as incident response or security architecture/engineering, which has strong relationships in place with lateral security teams. This organizational structure will help alleviate counterproductive posturing and politics that get in the way of results.

I can't stress enough that effective security is a cross-functional, cooperative effort. Walls between groups, whether caused by lack of communication, workflows that don't overlap, or big egos, need to be eliminated. When different security functions are siloed, critical information and intelligence doesn't get shared with the people who would benefit the most from it. My colleagues and I have seen many situations in which the security operations and incident response teams don't share "their" data with the threat intelligence team because they want to control where that data goes. This causes nothing but harm.

Other teams often think that threat intelligence is produced only from external sources by a dedicated team. They forget that the most important threat intelligence, and the first that should be generated and considered, comes from incident response teams who have visibility into what's happening within their own organizational infrastructure. That's what you should care about the most — not some report on a new exploit being used by some foreign threat against some other industry.

Tooling and Measurements

The final step in creating a successful strategic threat intelligence capability is defining the tools and workflows necessary to maximize the value of threat data.

Simplicity is the most important principle here. Indicators of attack or compromise (IOAs and IOCs), most commonly IP addresses, domains, and file hashes, are important for immediate response to ongoing attacks, but it's adversary TTP identification that is necessary for exposing risk.

Generally, a database is required to store and share data. Something as simple as sharing EverNote/OneNote notebooks may be sufficient. Threat intelligence teams should avoid the trap of overthinking these tooling and workflow requirements. Spending years designing a threat intelligence database system to store analyst notes or IOCs and IOAs is a poor investment for any business.

Once the threat intelligence team is generating relevant threat deltas (RTDs) to feed threat category risk (TCR) inputs and to better communicate with senior stakeholders via the risk model output values, the next step is creating strong relationships with lateral security teams for improved security controls. This requires agreement upon communication channels and intelligence formatting so everyone can obtain and use threat data.

If the incident response team works with tickets created in a system of record like Jira or ServiceNow, then the threat intelligence team should accommodate that existing workflow. Peer security groups like incident response, vulnerability management, fraud, threat hunting, and security engineering should regularly talk about recommendations for security control improvements that cover both technical aspects and policy decisions.

I previously discussed communicating risk to senior business leaders in terms of changes in the probability of loss. These changes are driven by RTDs, so it is important to track the quantity and quality of documented RTDs.

In fact, RTDs are the most important metric for the threat intelligence team. Measurements of RTDs inform the

frequency of security control improvements generated by partner teams. Remember that security control improvements improve threat-category risk model inputs, whereas RTDs may degrade them (until a security control improvement is made). In simpler terms, the right security control inputs will create a narrower and more accurate range for risk assessment, while relevant threat deltas will do the opposite.

Don't waste time splitting hairs over whether something is a metric, a key performance indicator (KPI), or an objective and key result (OKR). Decide on terms and definitions that are acceptable to the business, and then begin consistently measuring.

Mean time to detect (MTtD) and mean time to resolve (MTtR) are common metrics for incident response teams, and they can also be adapted for threat intelligence. Specifically, you can measure the mean time to surface (MTtS) and the mean time to assess (MTtA) new threat actor TTPs. These are valuable metrics to show progress over time.

There are two primary outlets for new TTPs — offensive scenario creation and internal telemetry hunting.

Creating new offensive scenarios to test existing security controls may require collaboration with a red team if your organization supports one. Translating TTP instances into a proprietary security control validation platform (like AttackIQ) will achieve a similar result.

The threat hunting team (or the hunting function within the security team) should also convert newly discovered TTP instances into search criteria to be deployed in a SIEM or other telemetry database(s) in order to surface previously undetected adversary activity inside the business network.

Chapter 12


The Intelligence to Risk (I2R) Pyramid and Additional Considerations for Strategic Threat Intelligence



Strategic intelligence is concerned with high-level threat issues, like the intent, capabilities, and targets of adversaries. It is meant to inform risk analysis and guide decision makers like executives and boards of directors.

In the summer of 2022, I sat down with a few of my Recorded Future colleagues (Dylan Davis, David Carver, and Harry Matias) to articulate our process for producing strategic intelligence. The result of our discussions was the Intelligence to Risk (I2R) pyramid, a framework that intelligence professionals and business leaders can use to increase transparency and cooperation between the two groups. The framework is a pyramid because each

successive layer is a smaller domain. As you move up the layers, performance becomes more difficult and ultimate action more challenging.



In the private sector, executives reading intelligence reports are often left wondering, “So what? Now what?” They are not given a way to connect specific threats with risks to the organization, nor are they given information to assess the degree to which alternative solutions could mitigate those risks. In short, they are unable to relate intelligence to actions and business outcomes using the language of business: risk.

Further, today’s CISOs are expected to present directly to their organization’s board of directors. The CISO’s ability to connect the dots between intelligence and recommended actions has a major effect on the board members’ confidence in those recommendations, and how likely they are to approve them.

The I2R pyramid outlines a framework for refining data into intelligence, and intelligence into recommended actions and a clear story about how those actions will reduce risk to the organization.

Let’s explore the various pyramid layers.

Layer 1: Events, patterns, anomalies

The base layer focuses on the building blocks of intelligence: events. Relevant events span an enormous range in both the physical and cyber worlds. From a shot fired on the Himalayan border between India and China, to a new implant framework open sourced on GitHub, to an attacker attempting

to bypass a web application firewall (WAF), billions of events transpire each day. They add up to a staggering amount of data to collect, parse, and analyze for intelligence.

Single events can be evidence of threats. In other cases, events can be grouped to reveal patterns that indicate attacks. Events that deviate from patterns (anomalies) can also signal malicious activity.

Your organization should have broad exposure to disparate events. Analysis starts with awareness. A big net is required. After you establish event awareness, you can see patterns and anomalies start to emerge from your analysis of those events.

Layer 2: Threat implication

The second layer, threat implication, is a critical step in the process because it requires a human brain (or extremely sophisticated AI). The process of extrapolating implications from events is best described as “second-order thinking.”⁵⁹ Sometimes implications are apparent, but it’s the less-obvious (unintuitive) ones that are most important for building toward a proper risk assessment.

In my experience, there is a wide spectrum of analyst abilities in this second-order thinking domain. It is helpful for analysts to have an inclination toward broad information consumption so they can correlate cybersecurity, business, and geopolitical data. However, I’ve noticed that the most talented analysts also have a native ability to think beyond the obvious and combine disparate data points into a cohesive narrative. The ability to regularly produce strategic intelligence hinges on a workforce capable of examining a very wide range of data and teasing out second-order threat implications.

Layer 3: Control validation

The third layer, control validation, determines whether an event, pattern, or anomaly (Layer 1) is a threat or a risk. A threat does not become a risk if existing controls are sufficient to mitigate it or the organization can transfer the risk to another party through a cyber insurance policy or some other means.

59. <https://economictimes.indiatimes.com/wealth/earn/what-is-second-order-thinking-and-how-you-can-use-it-to-succeed-in-your-career/articleshow/78587426.cms>

Maximum input from a broad cross-section of security professionals in an enterprise is ideal for accomplishing a thorough assessment. A security engineer may have detailed knowledge of controls that prevent BGP hijacking or denial-of-service conditions. A different group may have more-detailed knowledge of endpoint coverage, while another one may be able to speak to patching routines. An identity and access management (IAM) team may be able to evaluate authentication mechanisms. The broader the coalition used for control validation, the greater the likelihood of identifying control gaps that translate into risks for a business.

Layer 4: Recommendations

Layer 4, recommendations, falls between control validation and risks because properly contextualizing a risk with a recommendation for remediation helps decision makers think through the best course of action. Sometimes the action can be simply accepting a risk as part of doing business.

Business justification is a crucial part of a recommendation. Ideally, recommendations involve both short-term actions (less sustainable, but justified by the small investment of resources) and long-term actions (where higher resource investments are justified by greater benefits).

Layer 5: Upside risk and downside risk

Layer 5 involves describing and quantifying upside and downside risks. Upside risks are missed opportunities to achieve positive outcomes in areas such as speed to market, market share, customer churn, customer engagement, and innovation in new markets. Downside risks (as previously discussed) fall into five categories: legal or compliance failure, operational disruption, brand impairment, financial fraud, and competitive disadvantage.

This is the layer in which the security team should provide a quantitative or qualitative assessment along with a recommendation. Most executives are comfortable assessing a risk based on proximity and impact. Including one downside risk and one upside risk in a narrative about a threat improves the probability of successfully engaging executives.

Layer 6: Action

Finally, at the pinnacle of the pyramid, decision makers commit to actions. Reaching that point depends on the ability of the security team to convey a message that takes into account the business's current focus, goals, and constraints.

Optional: Executive feedback

An ideal I2R pyramid includes an additional layer allowing executives to provide recommendations about risk assessments that help the process evolve for the better. However, executive feedback is difficult to obtain regularly. Waiting for it can result in sacrificing the good while holding out for a perfect that may not materialize, so security leaders should be prepared to refine how the pyramid is used based on their own judgment.

I2R Pyramid Example: Ransomware

Let's see how the I2R pyramid process could work in the case of a ransomware attack.

Layer 1. An analyst observes changing patterns in ransomware techniques following the Colonial Pipeline attack. Specifically, ransomware operators are focusing their research and development efforts on Linux and virtual containers because of the growing likelihood that mission-critical data is being stored on those types of systems. Additionally, geographic targeting is moving away from North American enterprises toward other theaters, including government systems in less-developed nations. Theories about the causes of the change include diplomatic pressure applied to Russia following Colonial Pipeline to curb ransomware gangs in its territory⁶⁰ and the opportunity for higher payouts when government systems became inoperable.

Layer 2. Our analyst sees several threat implications that might affect the organization. First, subsidiaries in Europe, Asia Pacific, and the Middle East are likely to experience increased targeting. Second, the organization could be affected by attacks on its supply chain partners in those regions. Both of these threat implications are relatively obvious. A third, and

60. <https://www.politico.com/news/2022/01/14/russia-colonial-pipeline-arrest-527166>

perhaps less-obvious implication, involves the organization's reliance on government contracts. If certain government systems go offline for an extended period, revenue might be reduced significantly.⁶¹

Layer 3. When assessing control validation, the analyst works with other groups in IT to answer questions like:

- In the event of a ransomware attack, how much time is required to restore critical systems from backup?
- Do we have offline backups? How frequently are offline backups updated?
- Do we have cryptocurrency available in the event we decide to pay a future ransom? Will our insurance policy cover ransom payments and negotiation logistics?
- What is the efficacy of our internal phishing training program?
- How effective is our endpoint detection technology for identifying Cobalt Strike and similar precursor tools for ransomware payloads?

Layer 4. Depending on the answers to the Layer 3 questions, the analyst may consider recommendations such as:

- Short term, invest in additional backup capabilities
- Long term, invest in comprehensive hardware-based MFA
- Long term, build a purple team to simulate both attackers using ransomware TTPs and defenders using existing controls
- Long term, purchase or renew a cyber insurance policy with ransomware coverage

Layer 5. Our analyst articulates and quantifies the impact of risks including:

- Legal or compliance fines based on national or regional regulations
- Brand impairment leading to lost revenue

61. <https://restofworld.org/2022/cyberattack-costa-rica-citizens-hurting/>

- Increased costs of customer acquisition and retention
- Operational disruption and incident response and recovery costs

Layer 6. The final decision by management might involve acting on any or all of the recommendations. It might also include investing in intelligence resources and moving further left of “boom,” strengthening third-party exposure analytics, putting more effort into following in-the-wild exploitation of vulnerabilities, and assigning staff to detecting typosquatting and removing rogue domains before they can be used by ransomware gangs for phishing campaigns.

You can read more about the I2R pyramid and see additional examples of its application at: <https://go.recordedfuture.com/hubfs/reports/i2r-framework.pdf>.

Three As for Addressing New TTP Instances

After our walk through the I2R pyramid, here is a deeper dive into adversary TTP discovery and control validation (the first three layers of the pyramid).

Don’t assume that your security vendors are testing their products against real-life attacks and building in defenses. At Recorded Future, we regularly hear from CISOs who are using intelligence to identify control gaps in products, specifically EDR solutions.

To identify relevant threat deltas, you need an efficient and iterative workflow around TTP instances. It should focus on three phases (the “three As”):

- Awareness
- Assessment
- Amelioration

Information from the six sourcing buckets mentioned in Chapter 11 can help you maintain awareness of new TTP instances. However, you must have broad data access and smart alerting logic.

Assessment and amelioration depend on the knowledge and skills of your analysts. Once a potential new malicious tool or TTP instance is identified, dissection of the associated offensive methods and techniques begins. This should include an assessment of current security control responses. Often this assessment phase requires manual intervention to properly emulate a TTP chain or build and operate a tool.

Security Control Validation

Security control validation can be a very effective way to test whether new TTPs can penetrate existing security controls and increase risks for the enterprise. However, traditional, third-party red team and penetration testing engagements leave gaps. Hiring an external group to test security controls on an annual or even quarterly basis may satisfy compliance requirements, but it's insufficient to address the complexity and changes enterprises experience daily. Also, when penetration testers aren't rigorous about tracking and trying new TTPs, the exercise becomes no more than a test of whether the SOC or the blue (defense) team recognizes the penetration tester's favored techniques.

Companies like Qualys and Tenable provide software that constantly scans internal systems for technical vulnerabilities. Similarly, companies like AttackIQ provide software that programmatically tests security controls against the latest adversary TTPs. This software enables iterative "wargaming" that mimics the speed at which adversaries adapt to defenses, and immediately identifies gaps in security controls.

Security control validation platforms provide a valuable source of information on new TTP instances. Also, CISOs can use security control validation scores to tell a consistent story about changing risk to the board of directors. In addition to charting progress against a compliance framework, these scores chart operational security improvements (or deteriorations) over time in a reliable way.

Of course, any metric is prone to tampering,⁶² and these scores can be manipulated to tell a better story. Security teams can game any system to make scores look better (or worse), just

62. <https://www.bbc.com/news/uk-25022680>

like the chief of police or the mayor of a city may reclassify certain crimes from major to minor to create the appearance of a drop in major crimes. But you can minimize tampering if you set your risk-reduction goals from the start, choose consistent standards to measure the outcomes of any changes, and communicate those changes consistently to stakeholders.

Workflows and Outcomes

Workflows can boost, or hinder, both efficiency and effectiveness. You should examine your key workflows to make sure they are providing the outcomes you need. Also, because in 2023 organizations of all sizes are struggling with a dearth of qualified human resources, you should aggressively pursue opportunities to automate analyst workflows. Increased automation will not only increase the productivity of your analysts, it will also improve outcomes and the ability to communicate those outcomes.

The following list of threat intelligence workflows is ordered from easiest to automate to more complex and requiring more resources to automate:

1. Detection of brand and domain abuse and intellectual property leaks
2. Enrichment of indicator of attack and indicator of compromise (IOA/IOC) data for SecOps
3. Exposure analysis, particularly of technology stacks and third-party vendors and suppliers
4. Reporting
5. TTP instance identification and assessment
6. Risk quantification

Let's look at the reasons and opportunities for automating each of these workflows.

Detection of Brand Abuse and Intellectual Property Leaks

Phishing and domain abuse often coincide, but not always. Domain abuse is concerned with identifying attempts to spoof an organization's domains via typosquatting (the creation of

domains that are slightly different from those of well-known organizations and are often used for phishing attacks, scams, and the sale of counterfeit goods). A typosquat domain doesn't immediately correlate to malicious behavior but should be monitored and proactively removed when possible. Suitable open source tools like `dnstwist`⁶³ offer options for generating comprehensive domain permutations toward subsequent domain registration matching.

Code repositories and paste bins also require monitoring. I know from experience that developers enjoy maintaining code repositories on public resources for convenience. I've even observed developers backing up their entire hard drive daily to public code repositories. But code may contain private access keys and proprietary content. (Today, software is the most valuable asset of many companies.) It's generally a bad practice to sync proprietary code to publicly accessible code repositories on shared resources like BitBucket or GitHub. Monitoring for sensitive disclosures should extend beyond code repositories to the general web.

In both cases — domain typosquatting and IP leaks — the workflow is straightforward. Scanning domain registries and the web can be handled by machines, and then humans typically assess new results when they surface. The assessment and amelioration pieces are difficult to automate with complete fidelity.

Enrichment

Enrichment is the process of finding data related to vulnerabilities and threats and using it to provide context that helps security teams eliminate weaknesses and respond to attacks faster and more effectively. This context can include descriptions from vulnerability databases, IP addresses and domains used by threat actors, information about the history and capabilities of malware variants, and insights about the threat actor groups and their tools and methods.

For example, security vendors like Qualys, Tenable, and Rapid 7 provide varying levels of programmatic vulnerability assessments against discovered assets in the network. Those data results should then be combined with enriched threat data in

63. <https://github.com/elceef/dnstwist>

a system of record. Similarly, enriched threat data from third parties should be stored in a system of record. Proper system integration automates routine discovery and data collection tasks so analysts can concentrate on the strategic work of judging the severity of technical vulnerabilities and deciding how to address security issues at vendors and suppliers.

To obtain a complete, current picture of enterprise exposure, it is important to include data from a variety of sources, including passive and active telemetry and open and closed sources.

Enrichment processes can be extremely labor intensive, so they should be automated as much as possible. To save time, most enterprises outsource to system integrators (SIs) the construction of systems that collect enrichment data from multiple data sources and combine it in a master record like ServiceNow.

Exposure Analysis, Especially for Technology Stacks and Third-Party Vendors and Suppliers

Exposure analysis involves detecting vulnerabilities in assets (such as servers, endpoints, and security devices) and weaknesses in security controls, gathering contextual information about these vulnerabilities and weaknesses, and using that information to identify corrective measures and prioritize remediation. It begins with understanding assets and their relationships in real time. The challenge is managing the complexity of third parties and the constant adoption of new technologies.

Intelligence helps address this challenge. It can deliver data on vendors and suppliers, including evidence of past data breaches and existing vulnerabilities, and provide context to prioritize patches for vulnerabilities.

Reporting

I previously described why it is dangerous to make daily reporting the primary vehicle for communicating strategic threat intelligence. They increase awareness, but frequently fail to generate follow-up action and communication. In the

same way, quarterly metrics like “number of threat reports produced” rarely inspire executives to take action. Remember that the goal of cybersecurity is creating operational outcomes that can be measured and communicated.

The exception to this rule is when operational outcomes can be summarized to inform business decisions, typically those linked to a budget. If an enterprise cybersecurity program is already quantifying risk, then reporting is straightforward because it can focus on justifying spending to close high-lighted gaps in security controls.

Periodic reports tracking the number of new adversary tools and TTP instances identified, assessed, and ameliorated (in concert with adjacent security teams) will provide insight into the value of the threat intelligence team’s workflow and the benefits to the business in terms of risk reduction.

You can automate reporting workflows using security orchestration, automation and response (SOAR) products. One common use case is creating real-time dashboard reporting. Automation can also speed up the process of adding data and context to reports so they can be used for decision making.

TTP Instance Identification and Assessment

Strategic threat intelligence workflows involve identifying and assessing the latest iteration of TTP instances across risk categories. (I say TTP “instances,” because in a given year there are few adversary TTPs with new core tactics; most are slightly evolved instances of previously identified TTPs.) Most of a threat intelligence team’s time should be dedicated to the TTP instance workflow because this is where human brains are required and deliver the biggest return on investment.

For example, phishing, as a subcategory of social engineering, remains a primary method for initial unauthorized access. Tracking the evolution of phishing campaigns is necessary for most organizations that wish to measurably reduce risk.

In 2023, adversaries are using different attachment types to evade traditional email security controls, including embedded Microsoft Office macros, JavaScript, Visual Basic scripts, object linking and embedding (OLE) content, HTA (HTML

executable) files, and more. Many of these files fetch additional Base64-encoded scripts from external web servers.

Adversaries also use email links that route to spoofed websites with pass-through authentication credential capture features.

In the case of BEC, executives are targeted with a legitimate-sounding request that involves moving money to a purportedly legitimate recipient. No technology is required for a successful attack beyond the ability to successfully place a few paragraphs of text in a target's inbox.

Thus, social engineering in general, and phishing in particular, represent a considerable risk to most organizations, even those with robust technical and process controls in place. Strategic threat intelligence workflows involve identifying and assessing the latest iteration of TTP instances across risk categories.

Technology can improve the part of the workflow focused on identifying TTP instances, but the assessment piece can require time and very extensive activities.

Once new TTP instances are identified, they can feed red team scenarios, control validation software, and new internal threat hunting scenarios.

When a red team is available, new scenarios should be built using the latest TTP instance iterations. The blue team should be attempting to identify red team efforts. If a red team isn't available, security control scoring software like AttackIQ is an alternative way to build scenarios and test controls.

The remediation part of the TTP identification workflow often requires collaboration with a security engineering or architecture group, particularly when the gaps in security controls are large.

For example, in the case of phishing, if security controls prove insufficient against a specific offensive scenario, then the remediation part of the workflow may include deploying new controls like improved email gateway inspection, new Active Directory Group Policy monitoring, and ongoing analysis of quarantined attachments.

Risk Quantification

Previously, I discussed how to use the threat category risk (TCR) framework to quantify risks so your organization can better understand the value of operational security outcomes. You should set up a workflow to ensure that your risk quantification activities are carried out systematically and that no steps, including the creation and use of relevant threat deltas (RTDs), are missed.

Attribution

Before moving on, let's address the value of adversary attribution.


General adversary attribution can be helpful because motivation informs methodology. Knowing why someone is carrying out a cyberattack can help us better anticipate their targets and the means they will use to perform that attack. And knowing *who* is carrying out that attack helps us determine why. If you understand an adversary's motivation, you can better anticipate the TTPs they may use in the future, the level of resources they have available, how persistent their attacks might be, whether their attacks are targeted or untargeted, and so on.

However, more granular threat actor attribution (like name, address, picture, and so on) is irrelevant to the security needs of private sector organizations (although government security teams may need more detailed attribution).

For example, being able to attribute an unauthorized intrusion to the Chinese Ministry of State Security (MSS) is helpful context for TTP analysis, but there is no benefit to obtaining more specific information such as the name of the individual hacker working for the ministry or the address of the office where they sit.

Chapter 13

Operational Threat Intelligence and the Outcomes Matrix



What Is Operational Threat Intelligence, and Why Do I Need It?

Operational threat intelligence automates the collection and analysis of threat data that can be used to uncover and block ongoing cyberattacks and campaigns by threat actors. Whereas strategic threat intelligence is primarily focused on adversary TTPs, operational threat intelligence is concerned with processing indicators and artifacts associated with attacks, such as vulnerabilities, IP addresses, domains, uniform resource identifiers (URIs), and file hashes. Operational threat intelligence can also provide vulnerability enrichment and metadata about internal technology stacks

and third-party exposures, which may lead to operational outcomes in the form of detection and blocking rule sets (think YARA or Snort).

I often advise small and midsize businesses and enterprises with fewer resources to start their threat intelligence program with operational threat intelligence. Strategic threat intelligence programs are worthwhile but require substantial resources (including time) to execute properly. Strategic and operational threat intelligence are complementary, but if you only have the resources for one type, operational threat intelligence delivers great bang for the buck.

Operational intelligence can alert enterprises to previously undetected malicious activity, especially when threat data from outside the organization is correlated with internal telemetry obtained from a SIEM solution or an analytics tool. Even when correlation is not automated, organizations can start doing it simply by making sure that log data is visible and readily available to security teams. Once correlation activities are producing alerts, developing comprehensive SOAR workflows can reduce “noise” and false positives.

Sourcing valuable external indicators can be challenging: it’s costly if you are using vendors and time-consuming if you build your own infrastructure. Spending extra time digging into sourcing is always a valuable investment because it allows you to import only those indicators that will provide value.

I’ll note here that many people make a distinction between strategic, operational, and tactical threat intelligence (“tactical” sounds exciting — “tacticoool”). That is, “strategic” produces a product for executive consumption, “tactical” equates to analyst workflows, and “operational” covers processes that should be automated. For our purposes, it’s enough to understand threat intelligence through the two frameworks of strategic and operational.

Operational threat intelligence is enabled by machines and it should be programmatically applied to enhance existing security controls. For example, key indicators can be sent directly to firewalls, web proxies, or internal intrusion detection systems to produce alerts immediately. This type of integration is extremely valuable for small (one- or two-person) IT departments that need to automate their way to value.

Regardless of available resources for telemetry, the quickest path to value from operational threat intelligence is to funnel it directly to security controls such as firewalls, IDS/IPS, EDR, web proxies, and DNS RPZ. These can use the information to recognize and block malicious activity.

Large-scale telemetry correlation should only be considered once internal log collection is sufficient and sustainable. If the internal telemetry isn't available, or is only partially available, then the value of operational threat intelligence can't be realized.

Good and Bad Indicators

When you use operational threat intelligence, it is important to distinguish between good and bad indicators. In this context, “good” means helpful for identifying attacks and unlikely to generate false positives, and “bad” means useless for finding attacks and likely to generate false positives and cause other problems.

The rationale for acquiring malicious indicators is that adversaries reuse TTPs, including infrastructure. Some threat actors are careful never to reuse infrastructure, but in my experience, most of them, even at the nation-state level, are lazy. They reuse infrastructure, following the old adage, “If it ain't broke, don't fix it.”

Let's look at an example of a good indicator. Security researchers discovered that the Cobalt Strike tool used by adversaries contains a profilable secure socket layer (SSL) certificate. A server on the internet hosting a Cobalt Strike certificate has very likely been used in a cyberattack or will be used in one in the future. IP addresses and corresponding domains for such servers would be very good indicators of attack (IOAs). You would want to distribute them to network and host-based security tools so they could block all traffic from those servers.

An example of a bad indicator is the IP address corresponding to a botnet controller located on a shared hosting platform. On such a platform, 10,000 domains may resolve to a single shared server IP address. If you blocklist that shared server IP address (through DNS RPZ, a web proxy, a firewall, or something else), all the legitimate resources hosted on the same server will become unavailable to the enterprise, potentially

disrupting the work of some employees and causing panic among others who will believe the internet is broken. You will also cause your security analysts triaging security alerts to waste time on false positives because the internal traffic destined for the “rogue IP address” is actually legitimate traffic heading for other websites on the same shared server.

Domains provide much higher-fidelity signals than IP addresses for blocking botnet controllers and produce fewer false positives.

Another example of poor fidelity indicators are IP addresses and domains of websites that serve out malvertising. Adversaries inject rogue advertisements into advertising networks, and those ads redirect site visitors to a malicious exploit kit landing page on a server controlled by the adversary.⁶⁴ But it is useless to block a website where the rogue ad appeared, because that website is not malicious, and neither are the vast majority of ads displayed there. (However, it may be useful to block the website hosting the exploit kit landing page, if it can be identified.)

The value of indicators can also be evaluated based on periodicity. For how long is an indicator malicious? One hour? One day? One week? An indicator that is only good for an hour is not very useful.

Before jumping into an operational threat intelligence workflow, consider the goals and the measurements. The overarching goal must be to reduce risk, but again, the devil is in the details around how to measure and communicate the risk reduction.

Measurement

Operational threat intelligence workflows create a number of intuitive metrics.

1. For example, let’s say you have a workflow that:
2. Programmatically imports lists of breached user credentials (username, email address, and password)

64. [https://www.hsgac.senate.gov/wp-content/uploads/imo/media/doc/REPORT%20-%20Online%20Advertising%20&%20Hidden%20Hazards%20to%20Consumer%20Security%20&%20Date%20Privacy%20\(May%2015%202014\)1.pdf](https://www.hsgac.senate.gov/wp-content/uploads/imo/media/doc/REPORT%20-%20Online%20Advertising%20&%20Hidden%20Hazards%20to%20Consumer%20Security%20&%20Date%20Privacy%20(May%2015%202014)1.pdf)

3. Compares the breached credentials to existing user accounts in Active Directory
4. Determines which of your employees have been affected
5. Automatically resets the passwords of those employees

The number of employee account resets per period is a meaningful metric to report to the business because you are materially lowering the risk of data breaches.

Another example is removing or shutting down typosquatting domains. This is an operational outcome that can be measured and communicated. You might set up an operational intelligence workflow to:

1. Scan domain registry services and surface new domain registrations that are permutations of your domains
2. Enrich the domain listings with WHOIS data, nameserver identity, and SSL certificate data for each potentially malicious domain candidate
3. Add a subset of those domains and the related data to incident response tickets
4. Send the tickets to the legal department or a third-party domain takedown service

A third example is the importing of malicious IP addresses, file hashes, and domains for blocking actions. It's straightforward to track and report on the number of blocks. However, correlation for detection is less straightforward. The numbers of alerts triggered and triaged aren't meaningful because many of them may be false positives. Rather, it's the final outcomes that are important to measure: metrics like the number of identified infections or the number of security control changes made based on those newly identified infections.

A fourth example involves technology exposure analysis (vulnerability management). When importing operational threat intelligence data for vulnerability workflows, you might want to measure the number of:

- Vulnerabilities whose severity scores were altered based on programmatic enrichment
- Vulnerabilities enriched with evidence of exploitation in the wild that affect remote code execution (RCE) on internet-accessible systems
- Number of vulnerabilities and associated exploits identified before an official NVD/CVE identifier is publicly issued

Using these metrics to change risk scores in the TCR model will lead to changing monetary loss values. Even if you're not sold on the value of quantifying and monetizing risk, these are still meaningful intelligence metrics to communicate to senior stakeholders.

The Outcomes Matrix

An outcomes matrix, like the one shown on the next page, can help you clarify your thinking about desired outcomes related to specific security challenges and how to measure the efficacy of those outcomes. It helps you understand what kind of intelligence can create operational outcomes that can be measured and communicated.

The top row of the matrix lists high-priority security and privacy challenges. These should tie back to the risk impacts that are most significant for your organization.

Next comes the intelligence row. As I discussed in the section about the I2R pyramid, all intelligence is derived from an event, a pattern of events, or an anomaly. For each challenge, we need intelligence that will provide indicators related to one or more of the risk impacts. For example, in the potential breach notification column, a security team needs awareness of new phishing domains and domains impersonating a brand in other ways. Similarly, the team needs alerts about new mobile apps impersonating a brand and references to the organization in criminal forums and marketplaces.

The consumption row of the matrix shows how intelligence should be fed into security tools and processes. Because most security teams experience resource constraints, optimizing consumption is crucial to “doing more with less.” Common

consumption mechanisms include email triggers and API integrations into security, IT management, and reporting systems.

Problems	Potential Breach Notification	Infrastructure Exposure Identification	Physical Harm Avoidance	Quick Attack Remediation	Vulnerability Exploitation Avoidance	Fraud Avoidance	Credential Unauthorized Access	Supply Chain Liability	Security Control Efficacy
Intelligence 	• Phishing domain • Malicious apps • Code leaks • DW access • Advertising	• Internet inventory • Asset exposures	• Terrorist campaigns • Employee/asset threats • Travel risk	• IOX/OC context and environment • Infrastructure compromises	• Active exploitation • Pre-NVD • Pre-CVSS	• Stolen payment cards • Merchant breaches • Proxy/VPN use	• Stolen credentials/tokens	• Vendor/supplier exposure analytics	• Adversary prioritization • Internal threat packages • New 'tools' • TTPs
	• Email reporting • API	• API system integration • Email alerting	• Alerting • Geospatial monitoring • API system integration	• Browser extension • System of record integration	• Scanner integration • System of record integration	• API system integration • Manual reporting	• API for SOAR playbook	• System of record integration • Intelligence cards	• Red team scenarios • Hunting team scenarios
Consumption 	• Domain/social media/app shutdowns • Legal action	• Exposed asset remediation	• Site security • Business continuity response • Executive protection	• Quicker event verdicts • Faster incident triage • Detect/block control actions	• Patch prioritization	• Active cards flagged • Account takeover prevention	• Active Directory/Cloud account resets	• Vendor/supplier account auditing/enforcement	• Security control validation • Internal threat discovery • Trend identification
Outcomes 	• Mean time to remove • ROSI	• New assets discovered • ROSI	• Physical/operational system disruption • ROSI	• Correlated detection events • ROSI	• Patch escalation • ROSI	• Cost of fraud declined • Approved vs. declined transactions • ROSI	• Mean time to identify • ROSI	• Exposure identification • ROSI	• Mean time to assess • Mean time to deploy • ROSI
KPIs 	• NIST CSF: ID.AM-1-4 • Reduce breach probability • Risk reduction	• NIST CSF: ID.AM-1-4 • Reduce breach probability • Risk reduction	• NIST CSF: DE.CM-2-3 • Improve resilience • Regulatory compliance	• NIST CSF: DE.AE-2-3 • Improve resilience • Regulatory compliance	• NIST CSF: ID.RA-1-1 • PR.IP-12 • PCI DSS • Regulatory compliance	• PCI DSS • Regulatory Compliance • Improve brand equity	• NIST CSF: PR.AC-1-7 • Risk reduction • Regulatory compliance	• NIST CSF: ID.SC-1-5 • ID.RA-2-5 • DE.CM-4 • Regulatory compliance	• NIST CSF: ID.RA-2-5 • DE.CM-4 • Improve risk assessments
Risk Briefing 	• NIST CSF: DE.CM-5 • DE.CM-7 • Reputation management	• NIST CSF: DE.CM-5 • Reduce breach probability • Risk reduction	• NIST CSF: DE.CM-2-3 • Improve resilience • Regulatory compliance	• NIST CSF: DE.AE-2-3 • Improve resilience • Regulatory compliance	• NIST CSF: ID.RA-1-1 • PR.IP-12 • PCI DSS • Regulatory compliance	• PCI DSS • Regulatory Compliance • Improve brand equity	• NIST CSF: PR.AC-1-7 • Risk reduction • Regulatory compliance	• NIST CSF: ID.SC-1-5 • ID.RA-2-5 • DE.CM-4 • Regulatory compliance	• NIST CSF: ID.RA-2-5 • DE.CM-4 • Improve risk assessments

© Recorded Future®

The next row, outcomes, is the most important section of this matrix. What do you need to achieve? How is security improved and risk reduced through a particular intelligence consumption workflow? Returning to the brand protection example, once a rogue domain is identified, the desired operational outcome is a takedown request before users are phished or data is exfiltrated.

After achieving an outcome, how do we measure and communicate it? The KPIs (key performance indicators) row shows meaningful metrics for intelligence teams. “Mean time to remove” rogue domains is a worthwhile metric. Once a domain is identified with intelligence, the time required to remove it by working with registries varies with factors like geography, but the time measurement is consistent and valuable when communicated to the wider security team.

Return on security investment (ROSI) is defined as the annualized loss expectancy multiplied by the percentage of risk mitigated, less the cost of security divided by the cost of security.⁶⁵ Although the variables in this formula can be difficult to pin down with confidence, it is still a better metric for security than conventional ROI.

Finally, the risk briefing row helps relate operational success measurements to risk management themes for improved storytelling at the board level. It shows the alignment between outcomes and security and specific compliance frameworks such as the NIST Cybersecurity Framework (CSF) and the Payment Card Industry Data Security Standard (PCI DSS).

The outcomes matrix also helps articulate the value of using automation and engineering resources to enable intelligence consumption. In 2003, no organization could operate without a firewall. In 2023, operational intelligence warrants the same status for any serious security program. The barriers to entry for API connectivity to security and governance software systems (SIEM, SOAR, GRC, etc.) have never been lower, thanks to generative AI’s capability to produce code on demand.

65. <https://ccdcocoe.org/uploads/2018/10/Economics-of-cybersecurity.pdf>

Chapter 14

Five Critical Cybersecurity Functions



Security Starts with Preventing, Detecting, and Removing Unauthorized Access

Fundamentally, security is about answering a single question: How do I give the right people access to the right systems for the right amount of time, while keeping the wrong people out? Preventing and detecting unauthorized access is a key objective of cybersecurity.

Five core focus areas are critical to reducing the risk of remote, unauthorized access for a majority of organizations. In this second edition I replaced “web security” with “cloud security” in the list of security areas, primarily because mod-

ern browsers have created significant headaches for threat actors attempting to execute malicious code. Additionally, businesses continue to accelerate cloud adoption as part of larger digitalization strategies. This chapter will focus on exploring the following security categories:

1. Identity and access management
2. Vulnerability management (technology exposures)
3. Third- and fourth-party risk (relationship exposures)
4. Email security
5. Cloud security

These are the categories to prioritize, especially if resources are short.

The Problem of Identity and Access Management

The greatest problem in modern cybersecurity is identity. Today, the internet underpins global commerce, personal finance, media, and many more essential parts of the world's economies and cultures. In all of these areas, there is no trust without identity — and identity is extraordinarily difficult to verify on the internet. Circumventing online identity verification mechanisms over the past few decades has generated unimaginable wealth for threat actors.⁶⁶

The security industry has made some progress with multi-factor authentication (relying on something you are or something you have in addition to something you know). Unfortunately, in 2023 passwords alone are still widely used for authentication. Biometric validation — authentication based on something you are, like your fingerprint, retina scan, and so on — for application single sign-on (SSO) is a substantial improvement, but the software built on top of biometric authentication will always be vulnerable to tampering. There will always be unanticipated methods to bypass the next generation of authentication tools.

Interviews with experienced criminals, and analysis of

66. <https://resources.infosecinstitute.com/cybercrime-and-the-underground-market/#gref>

campaigns by state-sponsored actors, reveal credential reuse to be a preferred mechanism for gaining initial unauthorized access. All businesses are potential targets because the supply of stolen credentials seems to be limitless. Criminals can easily obtain access to vast troves of personally identifiable information in criminal marketplaces. Naturally, stolen credentials of employees with the highest levels of system access (think administrators authorized to access Active Directory Domain Controllers) represent the greatest threat to a business.


Examples of threat actors reusing stolen credentials to gain unauthorized access to corporate networks include:

- Visma, a Norwegian managed service provider, which was attacked by APT10, a nation-state sponsored threat actor, for the purposes of industrial surveillance (see the diagram below)
- Target, the American retailer, which suffered a substantial breach and data loss at the hands of criminals searching for opportunity

These were not isolated incidents against small and unprepared organizations.

Abusing customer services for fraud and financial gain is also a popular activity in the underground economy. Due to the massive availability of stolen credentials obtained via breached databases, it's never been easier for adversaries to execute credential reuse.⁶⁷ This dynamic is possible because people are lazy — it's just much easier and more convenient to use one set of credentials to access multiple online resources.

67. <https://www.recordedfuture.com/credential-stuffing-attacks/>



The TTPs APT10 used to breach Visma's systems. Source: <https://go.recordedfuture.com/hubfs/reports/cta-2019-0206.pdf>

IAM That IAM

Achieving comprehensive identity and access management (IAM) throughout an enterprise requires significant resources, including time, budget, and skilled security architecture and engineering groups.

The traditional challenge for CISOs is maintaining a patchwork of IAM solutions for various legacy systems and applications coupled with partial coverage for multi-factor solutions. The adoption of cloud and mobile technologies is rendering traditional network boundaries obsolete and putting further stress on the scalability of legacy IAM solutions. Cloud access security broker solutions have recently emerged in part to address the gap created when organizations attempt to extend IAM security policies to cloud resources.

Strategic threat intelligence is a force multiplier for the IAM function because it uncovers new adversary TTPs for bypassing authentication, authorization, tokenization, and other IAM functions.

For example, when pass-the-hash (PtH) attacks emerged in 2014, Microsoft issued remediative guidance suggesting actions that reduced risk.⁶⁸ But organizations that waited for Microsoft to issue formal recommendations were often too late. Rapid awareness and assessment were necessary to address architecture deficiencies prior to Microsoft's issuing of formal recommendations.

Similarly, obtaining organizational credentials from database breaches helps prevent customer account takeovers (ATO) and future credential reuse on the network.

Measuring and communicating the value of outcomes in IAM is relatively straightforward because senior stakeholders understand the basic concepts of identity and access.

68. <https://www.microsoft.com/en-us/download/details.aspx?id=36036>

The Role of Intelligence in Vulnerability Management

Organizations using vulnerability risk management (VRM) often struggle to properly identify vulnerability exposure and to apply patches and workarounds in a timely manner. Heterogenous and legacy environments for hardware and software, together with a lack of confidence about complete asset inventory, make vulnerability management an extremely resource-intensive activity.

Additionally, in enterprises with multiple lines of business, there is often a justifiable reluctance to patch when the possibility exists of prolonged outages of mission-critical systems. Business owners often prefer to accept the risk of unauthorized activity when the alternative is system downtime. Operational threat intelligence can play a major role in VRM by turbocharging patching efforts and clarifying the value of vulnerability remediation activities to skeptical business units.

Although most vulnerability management processes rely on the common vulnerability scoring system (CVSS),⁶⁹ vulnerability management teams need additional context beyond the CVSS base scores to increase severity scores when appropriate. That evidence may originate from internal asset data or threat intelligence articulation. As discussed in previous chapters, a good threat intelligence workflow automates the process of enriching data about vulnerabilities and feeding the information into a system of record like JIRA or ServiceNow.

Threat intelligence in vulnerability management should be measured by the number of pre-CVE⁷⁰ (common vulnerabilities and exposures) vulnerabilities surfaced, and how often severity ratings are changed based on evidence of remote code execution (RCE) that could potentially affect internet-facing systems.

69. <https://nvd.nist.gov/vuln-metrics/cvss>

70. <https://cve.mitre.org>

Third- and Fourth-Party Risk

Vulnerability management and third-party risk form one logical continuum. Both functions address surfacing and assessing exposures iteratively. While the VRM team addresses the potential for exposure in internal technology stacks, the governance, risk, and compliance (GRC) team focuses on external exposures via third parties, including potential exposures in their technology stacks.

It's understandable that adversaries see third-party relationships as natural avenues for exploitation. They can piggyback on pre-existing relationships between enterprises and their vendors, suppliers, and other trusted third parties. Digital supply chains continue to grow in scale and complexity, creating even more exposure for organizations that don't perform due diligence on their business partners or monitor them on an ongoing basis.

Examples abound of enterprises being compromised through trusted third parties. In late 2017, [24]7.ai, an online chat vendor, was compromised and personally identifiable information (PII) was lost from many national retailers that had technical integrations with [24]7.ai.⁷¹ In early 2018, MyFitnessPal (an Under Armour business unit) was attacked and PII subsequently exfiltrated.⁷² Universal Music Group and MyHeritage experienced similar victimization via vendor relationships.⁷³ Visma and Target, as previously discussed, are two more examples of exposures that caused wider damage to connected third parties.

A more recent example, from February 2023, was the ransomware attack on Dublin-based ION Trading Technologies Ltd.⁷⁴ ION, an important vendor for financial services companies, provides software that clears trading transactions. When the software became unavailable because of a ransomware attack,

71. <https://www.forbes.com/sites/leemathews/2018/04/09/hacked-chat-service-exposes-data-from-best-buy-sears-kmart-and-delta/#3c0e9d8f3055>

72. <https://fortune.com/2019/02/14/hacked-myfitnesspal-data-sale-dark-web-one-year-breach/>

73. <https://threatpost.com/honda-universal-music-group-expose-sensitive-data-in-misconfig-blunders/132451/>, <https://www.bloomberg.com/news/articles/2018-06-05/hack-of-dna-website-exposes-data-from-92-million-user-accounts>

74. <https://www.wsj.com/articles/cyberattack-on-ion-derivatives-unit-had-ripple-effects-on-financial-markets-11675979210>

banks had to scramble to manually track trade execution for end-of-day reporting. The list of cyberattacks that cause third- and fourth-party effects is long.

Managing third- and fourth-party (the vendors' vendors) exposure begins with tiering organizations by the level of access permitted to the primary enterprise. It isn't as critical to oversee the office supply company restocking printer paper as it is to monitor the online human resources and payroll service provider.

Organizations should create and maintain a list of third parties with access to customer or proprietary systems and data. The list should be segmented by the time and level of access required. In an enterprise this task is typically the responsibility of a GRC group. That group should use a system of record to track updates and changes to the status of third-party relationships, as well as their compliance with security policies such as patching known vulnerabilities.


You can accomplish even more by applying a risk-based intelligence philosophy to other aspects of third-party risk, and by integrating operational threat intelligence into the existing third-party system of record.

Let's say your organization has a relationship with Acme Financial Services (we performed a risk identification exercise with them way back in Chapter 2). Acme also provides online payment transaction services. While you have an interest in Acme's long-term economic viability, its current cybersecurity disposition is even more important. You should use threat intelligence and vulnerability scans to monitor:

- Exposed API keys in open and closed web sources
- Acme's use of website technology versions known to have vulnerabilities
- Previous breach disclosures
- Evidence of commodity network infections
- Past and current infrastructure misconfigurations
- Unattended domain typosquatting

When new events occur in the categories above, an audit of the third party in question may be warranted. The GRC group needs the authority to initiate audits. If those audits fail, then GRC needs the power to terminate the vendor or supplier relationship. Culturally, this can be a difficult recommendation, especially if the third party is integral to business operations, but leaving the relationship intact may increase financial loss in the event of a data breach.

In the past three years, third party exposures have become ever more important to risk management. Evolving language in contracts is a signal that enterprises are trying to transfer risk and create additional legal recourse in the event that a vendor or supplier suffers a disruption or breach.



A timeline of supply chain attacks, September 2021 to September 2022 (Source: Recorded Future)

Email Security

Access to the corporate email system is still a non-negotiable capability for most employees. Unfortunately, email remains a primary attack vector, primarily through social engineering. Phishing remains stubbornly effective. The best technical controls, like email security gateways, won't prevent a persistent adversary from successfully phishing employees, primarily because phishers continue to discover new methodologies for bypassing even the latest security gateway techniques like detailed content inspection, domain history, and sender policy framework (SPF).

That doesn't mean email security gateways are ineffective — quite the opposite. But even a 0.005% success rate for a phishing campaign can mean serious damage for the enterprise.

A few high-profile organizations whose attacks began with phishing include the Democratic National Committee, Sony Pictures, and Xoom.⁷⁵

Operational threat intelligence can play an important part in a multi-pronged strategy for email security. For example, threat intelligence can provide:

- IP addresses and domains of servers and bots on the internet associated with spam and phishing campaigns
- “Chatter” on dark web forums about planned phishing campaigns
- TTPs used in recent phishing attacks

Organizations can use this information to shut down some attacks immediately, for example, by blocking network traffic from external websites used in these attacks.

In addition, intelligence can help threat hunters build playbooks. The art of threat hunting is identifying patterns and anomalies in telemetry (log data) that are likely to be indicators of malicious activity. A friend of mine describes the practice as “dumpster diving” — there's a lot of trash to search through to discover something useful. Threat intelligence can give threat hunters insights into what indicators and artifacts on networks, servers, and endpoints reveal about the presence of attacks based on phishing.

Finally, generative AI presents challenges and opportunities for email security. AI creates an asymmetric problem for defenders by increasing the quality and velocity of social engineering attacks. On the other hand, AI can identify social engineering campaigns by analyzing the flow of inbound mail. It remains to be seen if it will produce a reasonable signal-to-noise ratio, but the potential is large. Security teams must maximize the defensive capabilities of AI to keep up with AI-powered threats.

75. <https://www.washingtonpost.com/news/politics/wp/2018/07/13/timeline-how-russian-agents-allegedly-hacked-the-dnc-and-clintons-campaign/>, <https://www.engadget.com/2014-12-10-sony-pictures-hack-the-whole-story.html> and <https://www.marketwatch.com/story/the-strange-case-of-a-money-transfer-firms-missing-millions-2015-01-07>

Cloud Security

The TeamTNT “Kangaroo” and “What Will Be” attacks illustrate a trend of increasingly sophisticated attacks targeting cloud technologies.⁷⁶ When these technologies — Docker containers, Kubernetes clusters, GitHub code repositories, hypervisors, and so on — are layered onto infrastructure, complexity results.

There are three primary security issues associated with mass-scale cloud adoption:

- A current and accurate attack surface view is difficult to maintain when digitization projects are moving quickly.
- When organizations own less of the computing infrastructure, they receive less-granular telemetry and lose visibility into activity on networks and systems.
- The opportunity for cloud instance misconfigurations is high, particularly with IAM permissions.

For cloud adoption, the security world is reconsidering the confidentiality, integrity, and availability (CIA) model in favor of the distributed, immutable, and ephemeral (DIE) model. This is a positive development. Thinking differently about security for a modern, cloud-based IT stack is beneficial.

Intelligence plays a primary role in updating attack surfaces through both passive and active infrastructure scanning, combined with vulnerability assessments. My conversations with CISOs suggest that the cloud access security broker (CASB) market still has plenty of room to develop more-satisfying products.

Addressing a lack of security telemetry begins with maximizing event logging collection via API. As organizations continue to move to cloud platforms and configurations and adopt SaaS applications on a large scale, improving log aggregation and analysis will only become more important in the future.

76. <https://blog.aquasec.com/new-malware-in-the-cloud-by-teamtnt>

Successful Business Is Risk Management

We've just about reached the end of the book, so let's review. Businesses (organizations of all kinds) require effective cybersecurity to safely operate information systems and maintain CIA (or DIE in the cloud). However, unlimited spending on security controls in perpetuity is obviously a non-starter. Enterprise risk management (ERM) is just that, management. To properly manage risk and plan investments appropriately, organizations need intelligence.

Building intelligence programs requires thoughtfulness around challenges and desired outcomes. Strategic intelligence involves connecting intelligence to business decisions through frameworks like the intelligence to risk (I2R) pyramid. Operational intelligence — through automated workflows — acts as a critical control that improves every security function. Both facets of intelligence are relevant to building internal coalitions for cyber risk management. An understanding of intelligence requirements and capabilities begins with prioritizing the five risk impacts that all cyber threats potentially cause: legal or compliance failure, operational disruption, brand impairment, financial fraud, and competitive disadvantage.

Communication about cyber risk may involve quantitative expressions, achieved through estimation and Monte Carlo simulations, or binary designations and storytelling, as most executives and boards of directors prefer (for the moment).

Finally, the world of security continues to rapidly change, particularly because of digitization and generative AI. To prevent adversaries from gaining a perpetual asymmetric advantage, organizations need intelligence to enhance and accelerate their security processes.

About the Author



Levi Gundert is Recorded Future's chief security officer, a role in which he leads the continuous effort to measurably decrease operational risk both internally and for clients. Levi has spent the past 20 years in both the public and private sectors, defending networks, arresting international criminals, and uncovering nation-state adversaries. Levi previously led senior information security functions across technology and financial enterprises. He is an author, a trusted risk advisor to Fortune 500 companies, and a prolific speaker, blogger, and columnist.

The Risk Business

Second Edition

What Leaders Need to Know About Intelligence and Risk-Based Security

What's the best approach to enterprise security? The industry has long focused on threat-based or compliance-based approaches — but many organizations struggle to balance technical tools and practical outcomes.

The answer is to focus on reducing risk.

In this book, Levi Gundert, chief security officer at Recorded Future, draws on his decades of experience to develop a comprehensive cybersecurity framework emphasizing risk over threats. The key, he argues, is for the technical side of any organization to present intelligence to executives and other decision-makers in terms they understand: the language of risk. All operational outcomes must be framed in a clear, concrete way that tells a story of profit, loss, and risk reduction.

The second edition of *The Risk Business* greatly expands on these themes with updated and new chapters on a taxonomy of the five leading types of risk impact (legal or compliance failure, operational disruption, brand impairment, financial fraud, and competitive disadvantage), the importance of second-order thinking for effective intelligence assessments, the Intelligence to Risk (I2R) Pyramid, and other conceptual frameworks that will give leaders a decision advantage in a complex, quickly shifting cyber threat landscape.

ISBN 978-1-948939-38-6



9 781948 939386 >