

FUSION CENTER THINKING

Extending Cyber Threat
Intelligence to Combat
Fraud, Geopolitical
Threats, Cybercrime,
IP Theft, and More

DAVID CARVER & LEVI GUNDERT

FUSION CENTER THINKING

Extending Cyber Threat Intelligence
to Combat Fraud, Geopolitical Threats,
Cybercrime, IP Theft, and More

BY DAVID CARVER & LEVI GUNDERT



CYBEREDGE
PRESS™

Fusion Center Thinking

*Extending Cyber Threat Intelligence to Combat Fraud,
Geopolitical Threats, Cybercrime, IP Theft, and More*

Published by:

CyberEdge Group, LLC

501 E. Las Olas Boulevard

Suite 300

Fort Lauderdale, FL 33301

(800) 327-8711

www.cyberedgegroup.com

Copyright © 2025, 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, 501 E. Las Olas Boulevard, Suite 300, Fort Lauderdale, FL 33301 or transmitted via email to info@cyberedgegroup.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 book and eBook for your organization, contact our sales department at 800-327-8711 or info@cyberedgegroup.com.

ISBN: 978-1-948939-49-2 (Paperback)

ISBN: 978-1-948939-50-8 (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:

Editor: Susan Shuttleworth

Graphic Designer: Colleen R. Abel

Production Coordinator: Jon Friedman

Foreword and Acknowledgements

As a result of Russia's invasion of Ukraine, my 2022 was dominated by clients reaching out to ask how to think systematically about cyber threats in the light of geopolitical fault lines.

As I drew up materials for reports and training sessions, I noticed a lack of clear and concise standards for evaluating cybersecurity with respect to non-cyber fields such as geopolitics. I told a colleague, Sam Lewis, that if we couldn't find a book analyzing non-cyber influences on cybersecurity, we should write one. She agreed, and the idea for this book was born. Sam is one of the deftest reviewers of geopolitical analysis I've ever met, and I want to thank her for saying "yes, and."

Candace Moix started turning the idea from a what-if concept into actual pages of actual writing. She also drew up the blueprints for the book's discussion of "What makes good intelligence?"

Levi Gundert, the book's co-author, provided commentary on the practical aspects of fusion center thinking, including our discussions of building teams, setting organizational standards, and aligning intelligence requirements with business risk. I recommend that you read his excellent book on cyber risk assessments, *The Risk Business*, available at <https://go.recordedfuture.com/the-risk-business-second-edition>.

Many other people contributed to this book. I'm especially grateful to three exceptional managers: Maggie McDaniel, Kevin Kelleher, and Seth Whitten. They gave me the latitude to write a book in between running finished intelligence projects.

The team I work with every day at Recorded Future keeps me honest. I deeply appreciate the passionate thoughtfulness of past and present team members in guiding me toward the right analysis for the best outcomes. (Go SUPER! go APEX! go ROR!)

If not for certain people, I'd be a much worse analyst, or not an analyst at all. My thanks to the hacktivist trackers at iSIGHT, the information ops hands at FireEye, and all the Futurists who have forced me to be intelligent when creating intelligence.

Finally, my wife Jonah let me spend many evenings writing about analytic frameworks and risk assessments. She is the reason I want to keep the world safe.

David Carver

Table of Contents

Foreword and Acknowledgements.....	iii
Introduction: Why Cyber Can't Fly Solo.....	1
Problem #1: You Can't Fight Cyberattacks Just by Knowing About Cyberattacks	1
Problem #2: No Single Intelligence Group Knows Enough.....	2
A Case Study: Geopolitics.....	3
More Domains Where Cyber and Non-cyber Challenges Collide	4
The Bottom Line.....	5
Fusion Centers and Fusion Center Thinking.....	6
Two Essential Principles	8
About This Book	9
Chapter I: Why We Need Fusion Center Thinking.....	11
Voices for Fusion Center Thinking	12
Success #1: Government Fusion Centers.....	13
Success #2: Fraud Fusion Centers	14
Case Study: Leveraging Geopolitical Context at Boeing.....	15
Fusion Center Thinking Starts at Home	16
Voices from the Field: Perspectives on Fusion Center Thinking	19
Before Diving In: Questions and Concerns.....	22
Chapter II: What Makes Good Intelligence for Fusion Centers?.....	25
What Makes Good Intelligence?	25
Objectivity.....	26
Thoroughness	28
Conciseness	32
The Intelligence Lifecycle Framework	32
More Frameworks, and a Caution.....	35
Chapter III: Fusion Center Staffing and Data Sources	37
Humans: The First, Best Data Collectors and Analysts.....	37
Four Major Types of Data Sources.....	39
Internal Cyber Data.....	40
External Cyber Data	42

Internal Non-cyber Data43

External Non-cyber Data44

Artificial Intelligence: Complement, not Substitute.....49

Chapter IV: Leadership, Communication, and Technology 53

 Leadership54

 Communication55

 Hiring for Fusion Center Thinking55

 Technology.....57

Chapter V: Working Examples – Fraud Fusion Centers..... 59

 Fraud Fusion Centers in Action 60

Chapter VI: Risk Assessment with Fusion Center Thinking..... 69

 Risk Assessment Versus Forecasting.....69

 Fusion Center Thinking and Risk Assessment72

 What to Ask from Fusion Center Assessments.....73

 Risk Assessment Frameworks for Fusion Center Thinking76

 A Final Example 81

Chapter VII: Critical Thinking within a Logical Framework..... 83

 Potholes and Gravel on the Road..... 84

 Logical Fallacies85

 Cognitive Biases..... 90

Epilogue: Fusion Center Rethinking..... 97

 A Quick Review.....97

Introduction:

Why Cyber Can't Fly Solo



Problem #1: You Can't Fight Cyberattacks Just by Knowing About Cyberattacks

It sounds paradoxical, but even perfect knowledge about attackers' tactics, techniques, and procedures (TTPs) is not enough to fight cyberattacks effectively.

Of course, your cybersecurity teams must be able to interpret Sysmon logs, security alerts, and indicators of compromise (IoCs) in order to detect compromised endpoints, social engineering tricks, privilege escalation, data exfiltration methods,

and all the other elements of attacks in play. But if your enterprise includes hundreds of thousands of devices, software workloads, data stores, user identities, and other information assets, you can't monitor and protect all of them at the same time. In fact, most security operations centers (SOCs) investigate only a fraction of the alerts generated by their current security systems.

However, if you obtain threat intelligence, and if you can perform cyber risk assessment based on knowledge of geopolitics, criminal organizations, human psychology, and your own business practices, you will be able to:

- Identify the adversaries most likely to attack your organization
- Understand the motivations and anticipate the TTPs of those adversaries
- Set priorities so you can counter attackers by protecting your most valuable assets with the right tools and resources for detection and response
- Make the right investments to prevent similar attacks in the future

Today, most cybersecurity groups understand these advantages and have analysts or teams dedicated to threat intelligence and cyber risk assessment. However...

Problem #2: No Single Intelligence Group Knows Enough

Modern private enterprises and government agencies face a broad spectrum of risks that cannot be fully anticipated or analyzed by any one intelligence analyst or group. Threats associated with international conflict, organized crime, espionage, digital transformation, AI innovation, and ideological fervor involve both cyber and non-cyber components. They can only be assessed, prevented, and mitigated through the collaborative efforts of cyber threat analysts and experts in other domains.

A Case Study: Geopolitics

In a survey of CEOs in late 2021,¹ consulting giant EY (formerly Ernst & Young) found that over half the respondents (55%) answered “yes” to the question, “Are geopolitical challenges forcing you to adjust strategic investment?” In a related report on geostrategic challenges to businesses, EY predicted that businesses would need to respond intelligently to a host of geopolitical problem areas, including worker mobility (a critical issue at the outset of the COVID-19 pandemic), protectionism (including the first Trump administration’s policies targeting China), and a worldwide reaction against globalization in supply chains.²

And CEOs faced these challenges before Russia’s invasion of Ukraine in February 2022. That invasion suddenly forced companies and governments across the globe to evaluate with whom and where they could do business, and whether not doing business might invite retaliation.

Resistance to the post-Cold War international order has been at work for the last decade and will continue to play out for many years. Longstanding traditional alliances have splintered or been reconfigured, and power centers once thought invulnerable have shown cracks. For example, in March 2024, Chinese President Xi Jinping invited executives from companies like Qualcomm and FedEx to Beijing to reassure them of the benefits for foreign investment in the country. Such a plea for international business, in its acknowledgement of weakness in the Chinese economy, would have been unthinkable even a few years prior. As we finish writing this, tariff policies in the second Trump administration have re-introduced protectionism as a driver of economic conditions and responses.

1 EY 2022 CEO Outlook Survey: <https://www.ey.com/content/dam/ey-unified-site/ey-com/en-gr/campaigns/ceo/documents/ey-ceo-survey-global-report-eurozone.pdf>.

2 EY 2022 Geostrategic Outlook: <https://www.ey.com/content/dam/ey-unified-site/ey-com/en-gl/insights/geostrategy/documents/ey-2022-geostrategic-outlook.pdf>.

Businesses as well as governments have faced the harsh effects of developments in information technology and security, including wholesale bans on technologies and companies from other nations. Inspired by China's [Great Firewall](#), in 2023 the Russian government [experimented](#) with disconnecting from the global internet in favor of a tightly controlled "Sovereign Internet." The West has adopted a hardline stance against the use of Chinese [hardware](#) and [software](#) - most notably efforts in late 2024 to ban the video-sharing platform [TikTok](#). Nations in and adjacent to the BRICS (Brazil, China, India, Russia, South Africa, Egypt, Ethiopia, Iran, Indonesia, and the United Arab Emirates) organization have explored monetary and technological ways to [overturn](#) the primacy of the dollar, such as using non-SWIFT payment protocols.

In its [Annual Report](#) on cyber threats for 2024, Recorded Future identified a host of ongoing and worsening threats associated with geopolitical flashpoints, including escalating cyber threat activity from China-linked groups targeting the US and Taiwan; an increase in nation-state abuse of generative AI for influence operations; and Russia-linked attacks (both online and physical) against critical infrastructure in Ukraine.

These events created significant threats that menaced business and government organizations across the globe. Those that could anticipate and navigate the complex, fractured, geopolitical landscape were usually able to avoid or minimize significant losses, and in some cases even find new opportunities for competitive advantage and growth.

More Domains Where Cyber and Non-cyber Challenges Collide

Geopolitics is only one example of a domain (in the sense of "a sphere of activity or knowledge") where cybersecurity and non-cyber challenges intersect in ways that can hinder – or help – business enterprises and governmental entities. Others include:

- Financial products and services
- E-commerce (particularly in segments vulnerable to fraud)

- Information technology
- Aerospace and defense
- Manufacturing and logistics
- Media and web services
- Any industry subject to government and industry regulations (because regulations can make unfortunate security lapses into fineable cases of security negligence overnight)

The Bottom Line

We can express the burden and opportunity of today's threat landscape like this: **If you want to be great at cybersecurity, you need to consume as much intelligence as possible about all the domains that affect you.**

Requirements to acquire and use threat intelligence effectively include:

- Familiarity with a range of domains relevant to your organization
- *Critical thinking skills* among producers and consumers of intelligence so they can make sense of everything that's happening in the world
- An effective way to *filter and interpret intelligence* amid a flood of raw data
- A *strong team* to interpret and make use of intelligence given the amount of noise in cybersecurity reporting
- The ability to *forecast and perform risk assessments carefully and accurately* in order to anticipate and prioritize emerging threats

In short, your organization needs the right intelligence and analysts who understand cyber threats *and the contexts* in which they arise. The required knowledge about geopolitical, criminal, and ideological adversaries and their motivations and targets extends well beyond technical expertise alone.

Fusion Centers and Fusion Center Thinking

But how do you assemble a team with knowledge about both cyber and non-cyber threats? Equally important, how do you organize that team so that participants with different backgrounds, experiences, and perspectives regarding physical and cyber security can work together effectively?

One proven model is the fusion center. A **fusion center** is an analytical hub with a team that can ingest, synthesize, and produce intelligence based on knowledge about cybersecurity and one or more additional domains such as geopolitics, fraud, crime, technology, finance, or climate.

The best examples to date fall into two categories:

- **Government fusion centers** that share threat intelligence across multiple national and local government entities, law enforcement agencies, first responder organizations, and others
- **Fraud fusion centers** that bring together fraud prevention and cyber threat intelligence (CTI) teams within a single organization (often financial services or ecommerce companies) to anticipate, detect, and prevent fraud

We will discuss these examples in the next chapter.

It is important to note that a fusion center does not need to be a single department in one location. It can include analysts and stakeholders in a virtual organization distributed across multiple locations. Further, a fusion center does not need to be confined to a single enterprise; its members can belong to multiple organizations within an industry or geographical area.

In fact, the most important factor for the success of a fusion center is not its administrative or physical structure, but rather what we call “fusion center thinking.” **Fusion center thinking**

is a view of security intelligence that considers cyber threats in light of the non-cyber triggers that cause them and the organizational components most likely to suffer from them.

Terms such as “convergent analysis” and “interdisciplinary analysis” are sometimes used to describe this approach. However, we think fusion center thinking better captures the analytical locus and purpose of this kind of intelligence framework because it implies a process for blending experience and knowledge from multiple domains.

Your organization can benefit from fusion center thinking even if you don't have a recognized fusion center. Individuals in management, strategy, finance, legal, marketing, and other disciplines can work with cybersecurity and threat intelligence teams to combine insights about changes, trends, and threats originating in both cyber and non-cyber domains.

When done right, fusion center thinking can help businesses anticipate threats much sooner and more accurately than a siloed intelligence program that concentrates only on “cyber data.” Fusion center thinking encourages analysts to investigate questions such as:

- Which threat actors and cyber spies are likely to be most interested in our business?
- Could a new or fast-developing technology create new vulnerabilities in our infrastructure?
- Is an impending war between two countries likely to affect our supply chain or cause new threat actors to target our organization for the first time?
- How could the evolution of cryptocurrencies affect the probability that our organization will suffer from ransomware campaigns?

The capabilities and benefits of fusion centers and fusion center thinking will vary by industry. The following table highlights some of the value fusion centers provide for a few industries.

Industry	Primary Threat(s)	Fusion Center Expertise	Fusion Center Value
Finance	Fraud	Cyber threat intelligence, fraud, economics	Reduce losses due to fraud
Software and Technology	Theft of trade secrets	CTI, geopolitics, IP protection	Protect IP and trade secrets
Manufacturing	Disruption of supply chains and manufacturing processes	CTI, operational technology, geopolitics	Maintain operations and protect revenue
Media	Social media account takeover	CTI, business analytics, social movements	Protect reputation

Two Essential Principles

There are two essential principles for creating threat intelligence that appropriately contextualizes cybersecurity with knowledge from non-cyber domains to support fusion center thinking:

- Everyone in an organization is an intelligence stakeholder: security intelligence should be seen as a shared project across an organization, not a narrow conversation between a small set of intelligence producers and consumers.
- A successful fusion center starts with fusion center thinking: the individuals who manage and participate in fusion centers need a guiding vision of how their collective information should satisfy the intelligence requirements of the entire organization.

About This Book

This book starts by describing the rationale for fusion centers based on the authors' own experiences and those of multiple cybersecurity personnel with whom we have worked. We then suggest how to select participants in the fusion center, collect the right data, build and run the center, and generate the desired output.

This book is for both analysts and business leaders. Although these roles are often very different, individuals in both need to understand the intersection of cyber threats with their larger context if they are to intelligently safeguard their organization from imminent risks.

We are confident that organizations that implement our guidance will be able to improve the efficiency and effectiveness of their security program. As we'll point out, access to more data is not inherently better, but an understanding of *what* to collect and *how* to interpret it will improve a business's responsiveness to risk, all the way from the tactical operations of security teams to the strategic vision of executives.

If the world is becoming a reflection of the internet, then very soon every sphere of activity and knowledge will intersect with cyberspace – and with it, cyberspace's need for cybersecurity. Businesses can try to ignore that evolution, but they cannot escape what Charles Dickens called the “unchallengeable reality” of the world as it is. Across geopolitical, economic, regulatory compliance, and criminology domains, from Silicon Valley to the South China Sea, for marketing managers, undercover agents, AI developers, and everyone else, there are thousands of connections to cybersecurity. We should be seeing them all.

Chapter I

Why We Need Fusion Center Thinking



Our argument is simple: We believe that no security program can consistently anticipate and prioritize cyber threats without some form of fusion center thinking.

It is very likely that your organization has already bought off on this proposition to some degree. For instance, if your security and threat intelligence teams are fighting advanced persistent threats (APTs), they are undoubtedly relying on non-cyber sources such as government and law enforcement announcements to help them understand state-supported and criminal adversaries and their motivations and targets. They have

probably developed a mindset that can fairly be called “interdisciplinary” because it encompasses knowledge of social and political factors in the outside world and the use of investigatory techniques (e.g., “follow the money”) well outside the traditional cyber technology sphere. Most likely, they acknowledge that ignoring non-cyber sources and rejecting interdisciplinary thinking would leave them hopelessly handicapped in their battle against APTs and other advanced threats.

Voices for Fusion Center Thinking

We are not alone in emphasizing the vital role of interdisciplinary approaches to cybersecurity and threat intelligence.

For example, Timo Steffens highlights the value of non-technical concepts like *cui bono* (“who benefits”) as core elements of APT attribution in his excellent book on the subject.³ Steffens has shown how much interdisciplinarity has become normalized for cyber intelligence. In an [interview](#) with Duo Security in 2022, talking about the resources that researchers now have when attempting to attribute threat campaigns to specific countries or groups, he comments on the frequent use of both technical and contextual material:

The most important development [in attribution over time] is that attribution usually does not start from scratch anymore. There is a large corpus of previous research that attributes certain malware families, or groups, or strategic patterns. Often, this is used to do short cuts: you link an attack to a group like APT28, and then you can just refer to previous attribution statements and can skip the later phases of attribution.

In a 2020 journal [article](#) titled, “Interdisciplinary Cybersecurity: Rethinking the Approach and the Process,” researchers at the University of Houston – Clear Lake show interdisciplinarity at work across cybersecurity domains such as criminology, economics, and law. They argue that “A holistic approach to cybersecurity is one that considers the many disciplines that produce cybersecurity professionals – technical and non-technical alike, in a coherent fashion.”

3 [Attribution of Advanced Persistent Threats.](#)

Finally, in his 2017 book *Cybercrime Through an Interdisciplinary Lens*, Thomas Holt talks about the blind spots that arise in the vision of analysts who fail to take multiple domains into account when thinking of cybercrime:

The gap in [computer science and the social sciences] regarding cybercrime has created siloed knowledge that limits its utility for policymakers and practitioners alike. Technology-focused cybersecurity researchers do not always take into account the humans who use technologies on a day-to-day basis. For example, an engineer could design a very secure system, but if it cannot be operated efficiently or effectively by humans, it will be unsuccessful. Social scientists do not always recognize or control for the technical flaws and system designs that enable offending or increase the potential ways that victim populations are vulnerable to compromise. As a result, they are unable to make policy recommendations that can fundamentally improve system-level cybersecurity policies and procedures.

Success #1: Government Fusion Centers

Many of the attributes of modern cyber threat intelligence were born out of the need for better intelligence collection and sharing after the September 11, 2001, terrorist attacks. Although the notion of fusion centers predates 9/11, their visibility and funding sharply increased in the aftermath, leading to the creation of 72 federally recognized state and urban area fusion centers under the aegis of the U.S. Department of Homeland Security (DHS). As described by a report from DHS in 2011, the goal of such centers was to share information to “maximize the ability to detect, prevent, investigate, and respond to criminal and terrorism activity.”⁴

The first generation of government fusion centers suffered growing pains that are highly instructive for organizations planning to create fusion centers today. Like many security

⁴ The report is available at https://www.dhs.gov/sites/default/files/publications/2011-national-network-fusion-centers-final-report_1.pdf

programs created in the wake of 9/11, they faced accusations of abusing civil liberties. Critics of the centers complained that their output was often irrelevant, opaque, and unnecessarily costly. Having been set up as part of a new layer of defense bureaucracy, such fusion centers were not easy to rethink or reorganize when their problems became apparent. And a “share everything” mentality sometimes led to violations of the principle of least privilege and the exposure of confidential data.

However, government fusion centers have matured, and today are supported in the United States by DHS (<https://www.dhs.gov/fusion-centers>) and the National Fusion Center Association (<https://nfcausa.org/>).⁵

Success #2: Fraud Fusion Centers

We have recently written about the need for fraud fusion centers, again pointing out the interplay between incident response, cyber threat intelligence, and fraud:

The fusion concept is vital in this fraud context because CTI analysts generally focus on identifying malicious fraud communities and associated TTPs (tools, tactics, procedures). Fraud analysts, conversely, are concerned with creating fraud analytics toward enhanced remediation. In traditional organizations, the CTI and fraud teams sit in separate departments and are constrained by politically driven data silos. Optimized information workflows for maximum fraud impact remain elusive. When implemented correctly, cyber-fraud fusion centers remove perverse incentives and enable a unified team to make the organization more effective.⁶

Fraud is an excellent starting place for organizations looking to eliminate intelligence silos. We highly recommend viewing fraud fusion centers as a model of how to get security intelligence producers to connect with areas of a business

5 You can see a brief video about government fusion centers in the United States at <https://nfcausa.org/wp-content/uploads/2020/09/video-2.mp4>.

6 The Need for Cyber Fraud Fusion Centers: <https://intelligence2risk.substack.com/p/the-need-for-cyber-fraud-fusion-centers>.

most likely to benefit from their intelligence. For example, while a fraud-focused fusion center is probably “table stakes” for a secure banking organization, a geopolitical fusion center might be necessary for a government agency, and a technology-focused fusion center could be ideal for a software development firm with proprietary designs and code.

Fusion centers can help teams from different disciplines share information with each other. In a white paper from the Payment Fraud Intelligence team at Recorded Future, the authors note that CTI teams likely have data that is useful to fraud teams but which the latter might not collect:

Unlike other datasets fraud teams work with, CTI inputs are largely external to existing fraud datasets. These external inputs provide fraud teams signals that can be crucial to getting ahead of the fraud kill chain. By moving ahead of the fraud kill chain, fraud teams can remediate fraud before financial losses, material costs, or reputational harm come into play. In this sense, one might formulate the underlying logic of effective CTI–fraud fusion as (internal inputs + external inputs) × analysis = fusion.

That formulation of a fusion center’s components works for any instance of fusion center analysis between cyber and non-cyber domains. *Internal cyber-related inputs* might be network traffic logs; *external cyber-related inputs* might be security vendor information on a certain country’s recurring APT activity; and external *non-cyber-related inputs* might be information about that country’s new need for information or technology related to a certain industry.

Case Study: Leveraging Geopolitical Context at Boeing

Dan Kropp, a former team lead of threat intelligence at Boeing, could not agree more that geopolitical context was crucial to his team’s setup for defenses against sophisticated nation-state groups.

One of us interviewed Kropp on the subject of how to translate knowledge of geopolitical incentives into strong intelligence

requirements. According to Kropp, a good approach to intelligence preparation for the battlefield is to model threats on both sides - for us, for them, and for our interactions. The Boeing team didn't focus solely on how nation-states used cyberattacks. Instead, they looked into nation-state threat groups and *all* of the types of adversarial actions they could take against the company's business.

After gaining a broad understanding of different countries, they adjusted their data collection requirements based on the threat group motivations and intentions that most closely overlapped with Boeing's business units. Through this exercise, some things became very obvious. For example, a country that has ambitions to project its military power globally would want access to space-based surveillance and communications and would be motivated to operate espionage campaigns looking for satellite technology.

Once Kropp's team talked to Boeing's engineers about which proprietary technology would be of most interest to nation-states, they built targeting packages for their red team. They also passed details to the company's incident response team, basically saying, "If you want to hunt for novel activity, this is where groups are going to target us." They did all of this based on relational mapping: What is the data a nation-state wants? Where is that data accessed (applications)? Where is it hosted (storage)? Where do the applications live (infrastructure)? Who is accessing it (user)?

Fusion Center Thinking Starts at Home

Before exploring interdisciplinary sharing between cybersecurity teams and experts in fields like geopolitics and the flow of technology across nations, we must emphasize that the most important source of non-cyber information for fusion center thinking is the business itself. Security analysts should not think about topics like cryptocurrency forks or People's Liberation Army troop movements until they have first studied the organization they protect. This means knowing about the business's primary competitors, its most valuable intellectual property, its leadership, its history, its reputation, its controversies, and its long-term strategy.

In a 2023 [article](#) for Dark Reading, Ross Haleliuk, head of product at cloud security company LimaCharlie, articulated multiple areas in which business knowledge drives effective cybersecurity prioritization:

Every day, security professionals are making decisions that impact their organization's security posture; they cannot rely on CISOs to be the only people with critical knowledge about the business. Understanding how the company generates revenue, how salespeople share information with one another and with their prospects, how finance teams access information when working remotely, and how vendors get paid is critical to properly securing the organization's environment. Statistically, it is more likely that a company will suffer a breach because of how some department has set up its business process, not because of the latest zero-day found by Apple (although learning about the latter might rightly be more exciting).

We interviewed Daniel Oliveira, an IT security lead analyst at a food and beverage company, about the value of understanding the business. He noted that employees at his company stay for quite some time. He has been with the organization for 22 years, and the person with the longest tenure has worked there for almost 40 years. Typically, employees move between different areas of the business and acquire a lot of knowledge about it. This helps when they are responding to an incident and doing any type of security architecture work. People with more than just an IT background add a lot of value. Oliveira said:

I started in a factory in Brazil, helping with the IT devices in the factory to support operations. This type of knowledge, when I moved to security, helped a lot because I know the impact if a system stops. I know what can hurt revenue if something stops.

Oliveira sees integration with other departments as an important value for the security team, especially since security teams typically have so little insight into the processes and needs of other parts of the business. He stated:

I think that we, as the IT community, should think more about business impact than IT impact. At the end of the day, if there's no company, there's no job for us. Sometimes IT people are not as concerned about issue eradication. Because I worked for a factory, I've seen how the manufacturing department takes issue eradication way more seriously than the IT department.

Integrating with other departments isn't just about understanding impact. It's also critical for gaining access to data about the IT environment that would be difficult or impossible to track otherwise. Furthermore, sometimes those other departments can handle decision-making for network assets that otherwise would have been relegated to IT by default.

Oliveira's team works closely with the supplier management team because, of course, every time somebody purchases a new system, it goes through them. They know they have to pay that bill, and if it's an IT expense, the security team needs to be in the loop to make sure that the expense is not for an implementation that is done by somebody outside the IT department. For example, during a recent holiday season, someone bought a system, hired a local developer, and put together a webpage to organize the company's end-of-year party. Unfortunately, they were collecting data on that website from employees at the company, but they never thought about privacy. So, if that website had gone live, it would have exposed employees' dietary needs and personal data. Sometimes people don't think about these types of issues. Oliveira commented:

Another department that doesn't come to most IT people's minds is trademarks. [T]he people responsible for trademarks have a lot of information because they're responsible for registered domains... Sometimes we noticed in our DMARC journey, when we started working closer with trademarks, that they had a lot of info that is quite useful for IT folks. And then because of that, we were able to streamline new domain services, so every time someone put together a new domain, it had to be approved by trademark, and then they would call the IT department and let them know, so that, for example, HR couldn't buy a system and not notify IT.

Oliveira also noted that sometimes his team finds a website associated with a .ru or .ua top level domain that has their company's name in it. A green analyst might look at that domain and say, "We need a takedown." However, the trademark experts have the data to tell the analyst that those domains are legitimately associated with the company's brand and should not be taken down. The company has a lot of lookalike domains, and instead of his team having to investigate each one and decide whether to take action, the trademark group can make the call.

Jaime López Sánchez, a delivery manager in a global security operations center, agrees with the idea that security departments should know much more about their overall businesses than they currently do.

In an interview, he conveyed his opinion that the main department that cybersecurity analysts should have more information about is corporate security, especially in industries such as energy and finance. In these big industries, corporate security departments are often staffed by people with law enforcement backgrounds (for example, ex-FBI employees) who like to work on their own and are reluctant to share information with other departments. López Sánchez has encountered this dynamic, but he has also seen a lot of success stories when the cybersecurity and corporate security teams work together.

López Sánchez also pointed out the benefits of collaborating with the teams responsible for vetting acquisitions and suppliers. These teams often have a good grasp of important changes in business strategy. Knowledge of these strategy changes helps the cybersecurity department detect brand abuse and treat new domains appropriately. Sometimes cybersecurity is the last department to know about big business decisions, López Sánchez said, but with regular collaboration, other groups can better support cybersecurity teams and vice versa.

Voices from the Field: Perspectives on Fusion Center Thinking

As part of researching this book, we talked to many individuals with experience in cybersecurity domains, such as cybercriminal tracking, security architecture, security program management, and strategic forecasting, about how informa-

tion from non-cyber domains has improved their ability to properly articulate risks and threats.

One insight came up again and again: it is critical to understand non-cyber domain knowledge *in the context of the organization the security team is protecting*. The most important non-cyber domain differs widely by industry. For an aerospace company working on cutting-edge satellite research, foreign espionage is a critical threat to understand. For a packaging company, operational downtime at a warehouse has a crucial impact. For consultants working with Fortune 500 companies, an understanding of larger security concerns provides some of the biggest wins.

Samuel Curet, who co-founded the company Zero Trafficking and is currently a senior investigator at SpyCloud, has spent years studying the interplay between physical criminal activity and cybersecurity. Having spent much of his early analyst career hunting the gangs behind human trafficking, Curet is very familiar with how these gangs rely on the internet:

Sex trafficking and drug cartels have worked together and operated as separate lines of business for the same group of people. But because prostitution became an online thing, it almost became a nexus point to all kinds of other bad activity that are not related to physical crime (such as romance scams, pig butchering, and digital brothels).

Curet noted that different types of cyber crime should be considered in light of whether a crime is completely online, or is simply a cyber layer on top of physical crime. These insights can only happen with years of experience in the field:

You have to look at the regionalization of cyber crime actors. You have people that have been manipulated for fraud scams in Southeast Asia, but then you have a cottage industry [for fraud] in Nigeria, which is a more permissive environment, and then you have places in Eastern Europe and Russia where, as long as you don't attack CIS [Commonwealth of Independent States] targets, you can do what you want. However, in all the time I've spent looking at human trafficking, there is something unique about it, and it's unique

because it exists in the physical world, and [is] enabled through online advertisements and digital personas. It really is a nexus point between the cyber crime space, the fraud space, and even to some extent the drug space.

Insikt Group is Recorded Future's threat research division. The leader of the group's cybercrime intelligence team agrees with the need for attention to detail when discussing cybercrime, as well as the need to be careful when using this term as a catchall for any illegal activity related to the internet:

Cybercriminal marketplaces and sources like Silk Road and AlphaBay showed where the physical world touches the cyber world. But the word 'cyber' can be ambiguous. Think of shops where people sell payment card data that was stolen via physical skimmers or a Magecart attack - do we classify both of those as 'cyber' threats just because both can result in stolen credentials being sold online?

As we have mentioned, another huge area of study that affects cybersecurity research is geopolitics. When we spoke to Kelli Vanderlee, senior manager of strategic analysis at Google Threat Intelligence Group, she confirmed that her background in Middle East studies was invaluable for spotting suspicious behavior in purportedly hacktivist activity:

If you remember "Bozkurt Hackers," this was allegedly a hacktivist group identifying itself as ideologically aligned with the "Grey Wolves" Turkish nationalist militant organization. They leaked a bunch of data from a Qatari bank in 2016, alongside political messaging critical of the Qatari royal family and Al-Jazeera, and also revealed that a famous Egyptian Muslim Brotherhood scholar, Yusuf Al-Qaradawi, was living in Qatar.

This messaging didn't make any sense - Turkey and Qatar both supported the Muslim Brotherhood and had been collaborating in the context of the Syrian conflict. That ideological inconsistency was the first big sign to me that something very odd was happening, and we should keep digging. Tracking this false

hactivist persona's behavior allowed us to recognize the same kind of activity happening with Anonymous Poland and subsequently Guccifer 2.0 and DC Leaks [which were ultimately tied to operatives of Russian intelligence agencies].

Before Diving In: Questions and Concerns

We've established that individuals involved in academia and cybersecurity communications recommend fusion center thinking, and we've shown that people across industries and roles also believe it to be a boon for overall organizational security. That said, we can imagine a few questions or criticisms of the idea, which we'll briefly touch on here before constructing the larger framework of fusion center thinking best practices.

Question: It's fine to talk about adding non-cyber domain knowledge to cyber threat intelligence, but my organization already has a hard time dealing with cyber knowledge by itself. How am I supposed to layer the rest of the world on top of that?

Answer: In Chapter III we talk about how to prioritize data collection with fusion center thinking in mind. Until then, our position is that properly conducting this type of intelligence should, over the long term, create less work, since the goal is to contextualize your threat landscape in such a way that you can rank the most common and most severe threats based on your environment, and focus primarily on addressing those.

Question: There are thousands of different aspects to my organization and environment that I could see contributing to a threat landscape. How do I choose what to focus on?

Answer: In Chapter VI, we'll look at the appropriate expectations for risk assessments. One of those is that a risk assessment should be based on a comprehensive and ranked audit of assets. If you know which of your assets are most important, you can narrow your field of vision to environmental factors that are most likely to drive threats to those assets.

Question: I'm an IT leader and I struggle with getting our users to take basic security precautions or take security training. How on earth would I get everyone in the organization to be an intelligence stakeholder?

Answer: Everyone in an organization already is an intelligence stakeholder, whether they're aware of it or not. The way we think about that should run in two directions: individuals outside of security should take security seriously, while those inside security should understand that they need to protect the intelligence requirements and daily tasking of their colleagues. So, even if you can't get the first thing to happen, you can certainly improve awareness of the second.

Question: I'm a producer and consumer of internal cyber threat intelligence, and I don't feel that other consumers and producers genuinely understand what I need for our relationship to be successful - it's as if we're talking past each other. Is fusion center thinking likely to change this?

Answer: As in any communication, if you're not using a shared language, you're not going to have a real conversation. Properly implemented fusion center thinking should bring consumers and producers of intelligence into better alignment by establishing common baselines of critical reasoning (Chapter II), common asset priorities (Chapter III), organizational buy-in (Chapter IV), and shared expectations of assessments (Chapter VI). We are comfortable assuming that, if you're having difficulty with your counterpart, it's because one of these foundations of fusion center thinking hasn't been set up.

Chapter II

What Makes Good Intelligence for Fusion Centers?



Before we describe how to select the right people to promote fusion center thinking in an organization, let's examine what attributes define good intelligence in this context.

What Makes Good Intelligence?

Intelligence should always be specific and trustworthy. Broad, ambiguous, or dubious claims are unlikely to be actionable. And if it's not actionable, what is the purpose of intelligence?

All intelligence output should be:

- Objective: as free of personal opinion and bias as possible
- Thorough: backed by enough research and analysis to answer any follow-up questions a reader might have
- Concise: as simple and straightforward as possible, consistent with the complexity of the issues

Objectivity: “Would People from Different Backgrounds Agree with This Assessment?”

Any intelligence analysis - whether about financial, cyber, or national security issues - should strive for objectivity as its core attribute. Without objectivity, analytic endeavors are unlikely to yield meaningful results that reflect real-world risk.

The greatest enemies of objectivity in intelligence assessments are biases. They include both cognitive biases and other types resulting from cultural, political, and religious backgrounds. Producers of fusion center intelligence must ensure that their assessments are not unduly colored by personal ideology, opinion, or preference.

While that *sounds* easy, in practice it is common to encounter statements of “fact” that are, in reality, highly suspect. For example, in intelligence reporting on Chinese government goals with respect to Southeast Asia, we have read dozens of assessments that attribute every suspected Chinese espionage activity to “support for the Belt and Road Initiative.” That attribution represents both a cognitive bias (analysts remembering the most prominent Chinese government initiative and ignoring alternative explanations) and a cultural bias (the idea that if the Chinese government has announced a public priority, it must be spying on its neighbors to support that priority).

There is no inherent problem if a team’s analysis diverges from conventional wisdom within a field, but it is critical to understand *why* it’s unique. Is an original assessment the result of novel data sets, or improved methodologies, or weaknesses in the research design and analysis of other analysts? Unique intelligence is desirable because it yields new insights and

potential solutions to complex problems, but only as long as the analysis is trustworthy.

Different intelligence disciplines inherently involve different levels of objectivity. For example, signals intelligence, or SIGINT, is typically fact-based technical reporting with little room for data interpretation. In contrast, human intelligence, or HUMINT, necessitates qualitative and subjective interpretation of human behaviors.

Fusion center intelligence assessments should take account of evidence from multiple disciplines, but give more weight to those with higher levels of objectivity. For example, cyber threat intelligence, typically focusing on technical data points, tends to be fairly objective. Geopolitical intelligence entails a degree of subjectivity in assessing foreign nation-state motivations, so conclusions in that sphere need to be examined more carefully.

A word about diversity in intelligence teams

For at least the last five years, managers hiring for intelligence groups have been bombarded with statements along the lines of: “Diversity is necessary for a strong intelligence team because different perspectives help analysts to avoid bias.”

We agree with this statement wholeheartedly.

However, when considered in a vacuum, it can be dangerous and inaccurate:

- Dangerous, because it can be taken to imply that the only justification for diversity is reducing bias
- Inaccurate, because it implies that diverse hiring practices by themselves can eliminate bias and produce more-accurate intelligence, when multiple [studies](#) have shown that having a diverse workforce improves performance *only* when the values of diversity are championed by leadership and accepted and understood by the rank and file

It is not a straightforward project (and certainly not one that everyone agrees on) to assemble a diverse team of intelligence producers based on conventional markers of “diversity” such as a balance of gender, race, and ethnicity that matches the

general population. But in our experience, increasing the level of diversity in analysts' backgrounds, expertise, and attributes will reduce stereotyping and bias and result in more-objective framing of the data inputs, especially when these values are encouraged by management.

Thoroughness: “Does This Intelligence Include Everything It Should?”

Good intelligence is based on good research; great intelligence is based on checkable research. To be described as thorough, an intelligence report must be the product of all relevant data and methodologies, and those data and methodologies should be fully available to the final reader of the report.

A thorough intelligence report does not need to be long. Rather, it should have just enough detail so intelligence consumers will understand its assessments and be able to check them against the underlying research and analytical framework.

For example, imagine that a business leader asks a fusion center to describe the likely cybersecurity implications of a major Asian power's incoming administration. To answer that question thoroughly, the fusion center analysts need to collect data on the cybersecurity implications of any political change, the most important or characteristic *political* elements of the new administration, and the most relevant *cybersecurity* factors to be affected by those political elements.

Once that initial research is done, the analysts might make an assessment along the lines of, “It is unlikely that the incoming administration's actions over the next two years will result in major changes to the cyber regulatory environment or to the targeting patterns of threat actors that we monitor.”

Some intelligence consumers will be happy just to hear an assessment from a team they trust: “Sounds good, close the book.” However, and particularly for intelligence that is driving costly or controversial decisions, other consumers will want to understand *why* the analysts came to certain conclusions.

Is the assessment based on reliable evidence? Is that evidence relevant? Is the logic leading from the data to the assessment valid? Are there immediate follow-up questions (such as, “What would have to change to result in significantly different outcomes?”) that aren’t answered in the report?

Effective information sharing and intelligence production depend as much on transparency as on high-quality inputs. If intelligence stakeholders cannot easily identify the data and logic that led to the final assessment, they may not trust it as a basis for an important decision or recommended behavior. Even worse, if an incomplete piece of intelligence leads to a wasteful or harmful action, the entire output of the intelligence producers may be called into question. Certain organizations may have to limit transparency to protect proprietary methods or sources, but this should be the exception to the rule.

Here, for reference, are bad and better answers to two intelligence questions that illustrate the importance of thoroughness:

- Question: Is it likely we will be affected by a ransomware attack in the near future?
 - Bad answer: “I said that a ransomware attack would likely affect us soon because it seems like this type of activity has been going on for a while.”
 - Better answer: “I said that a ransomware attack would likely affect us soon because (1) it’s the most common form of known threat against our industry in the past year, and (2) three of our industry peers, with whom we are regularly named as a set in the news, have reported attempts against their network from this ransomware group. I’d give it an 80% chance of an attempted exploit in the next two months.”
- Question: Do we need to take measures to protect ourselves from that North Korean hacker group that targets cryptocurrencies?
 - Bad answer: “There’s no way North Koreans would target us. We don’t do anything with crypto.”

- Better answer: “As an enterprise, we are probably not the priority targets for this North Korean threat group because they more commonly target individuals. However, this group does narrowly target software developers or other people likely involved in technical work, so we should consider security measures for employees matching those attributes.”

Analysts who are asked to further explain or justify their assessments should either have ready answers or know how to get them. If pressed for extra details they should be able to clearly show how those details contributed to the more straightforward final assessment.

Rigorous, thorough intelligence should meet the following criteria:

- The analysts considered all relevant data and hypotheses.
- The research design and methods are straightforward and transparent enough so they can be evaluated by external experts.
- The intelligence is replicable by other analysts who do have access to the relevant data (if it is feasible to give them access).

A word about estimative language

A final point about thoroughness, which is sometimes difficult for people outside the intelligence field to appreciate, is that the strength of intelligence reports must be assessed in two ways:

1. According to the trustworthiness of the sources of data and information (a *confidence* assessment)
2. According to how effectively the research can be used to explain or predict outcomes (a *likelihood* assessment)

For example, suppose that executives are thinking of acquiring a company in Finland. They're not sure what security threats such a move would expose the business to, especially since the Finnish firm's CEO has very vocally opposed Russian political interference. They ask a fusion center for an assessment.

There are two ways the fusion center team can respond to this kind of question.

The first way is to include in the assessment confidence and likelihood language aligned with a standard such as the one described in the [US Intelligence Community Directive 203](#), which is depicted below.




Figure 2-1: Probability language included in US Intelligence Community Directive 203.

The wording might be something like: “We assess with moderate confidence that acquiring this subsidiary will likely expose our company to increased risk of hacktivist activity from pro-Russian threat actors.” This is a perfectly serviceable assessment, but for readers who don’t know how to interpret the terms “moderate confidence” and “likely” it will be boring or ambiguous.

The second approach is to strip down the confidence and likelihood phrasing and clarify the connection between the fusion center’s research and the decision under consideration. Such a version might read: “There is a high chance that acquiring this subsidiary will spur Russian hacktivists to try to disrupt people’s ability to visit our website.”

The debate over whether to use the first or second approach has been ongoing in intelligence production circles for decades. Which approach organizations choose for the *ultimate* version of intelligence will vary based on final audiences, but all great analysts should be comfortable presenting assessments with the technical specificity of the first approach.

Regardless of the approach selected, fusion center analysts should be explicit about how much they trust their sources and how much confidence they have that their data explains or predicts a scenario. Also, they must exercise caution when they are not familiar with the trustworthiness of sources. For example, if analysts are asked to determine whether a physical protest will spur cyber threat activity but have not yet developed a detailed understanding of the protest’s causes and associated

cyberattacks, they are unlikely to know which sources of information are most trustworthy and what data from those sources is most reliable. Under those circumstances, instead of faking a confidence they do not have, the analysts should either consult with an expert in the relevant domains or assign lower confidence levels to the unfamiliar sources.

Conciseness: “Will It Be Understood and Lead to Action?”

Good intelligence is straightforward, simple, and concise. Verbose or jargon-filled reports will not be fully understood or processed, and will not lead to appropriate action. In fact, they usually lead to no action at all.

Sometimes, being succinct requires breaking habits learned in academia, where long explanations of methodology are often expected and encouraged. This doesn't mean analysts should shy away from complex problems or data. But when complexity is unavoidable, analysts should employ methodologies and structured techniques that make data visible and reasoning clear to non-technical intelligence consumers. For example, pictures, diagrams, flow charts, bulleted lists, and tables can be powerful aids to conciseness and clarity. Analogies and brief anecdotes can have more impact than lengthy paragraphs.

As an exercise, imagine you have only five minutes of a decision maker's time. How are you going to use it?

The Intelligence Lifecycle Framework

Analytical frameworks are useful for fusion center thinking because they help counteract cognitive and personal biases and the mind's tendency to process information in ways that require the least time and energy. They help shake the brain out of comfortable ruts. Some analytical frameworks are better than others for certain types of questions, but simply having a framework forces analysts to evaluate data in a deliberate, thoughtful manner without taking dangerous shortcuts.

Probably the most important framework for fusion center analysts to be aware of is the intelligence lifecycle. It has been a standard framework for intelligence production for at least 100 years, **if not longer**. Although imperfect, it has obvious utility for many situations.

The picture of the intelligence lifecycle below should be familiar to everyone who has worked in intelligence. But even experienced analysts should review each stage of the lifecycle in the context of fusion center thinking, since interdisciplinary analysis for business relevance places new demands on the model.




Figure 2-2: The intelligence lifecycle.

In **Direction/Discovery**, the primary question for a fusion center is, “**Why do we need this intelligence?**” The questions asked at the beginning of the intelligence lifecycle should be directed toward producing intelligence that is relevant to clearly defined needs and leads to appropriate actions. In the context of fusion center thinking, business units should clearly link their questions about geopolitical, technological, financial, and social impacts on cybersecurity to decisions they need to make, and they should inform the intelligence producers about these decisions.

In **Collection**, the primary question is, “**Where are we getting the data for this intelligence?**” We will cover this question in Chapter III. A challenge for fusion center thinking is deciding how much data in how many different topic areas is appropriate to collect.

In **Processing**, the primary question is, “**How are we processing the data that we collect?**” Deciding how to process data requires a combination of critical reasoning and the input of subject matter experts. Fortunately, this is the most mechanical of the stages in the lifecycle, and approaches that work for processing other types of intelligence don’t require much modification for fusion center thinking.

In **Analysis/Production**, the primary question is, “**How are we analyzing the data that we collect?**” Once more, critical reasoning and subject matter expertise are necessary for doing this well. Also, fusion center thinking must wrestle with how explanative factors from different topic areas influence each other. For example, in certain scenarios it might be worthwhile to examine how the availability of certain cryptocurrency laundering tools might amplify the volume of ransomware attacks. Our commentary above on confidence and likelihood assessments relates directly to this point.

In **Dissemination**, the primary questions are, “**Who is this intelligence going to? Does it help with a decision, and can the audience make that decision?**” This stage needs to be aligned with the questions asked in the Direction/Discovery stage of the lifecycle. Fusion center analysis should be created for and directed toward individuals who can make decisions based on its findings. For the Dissemination/Action stage to be successful, intelligence producers need to have previously determined the appropriate consumers.

In **Feedback**, the primary question is, “**Did producers create what consumers needed?**” A feedback loop must involve responses that go beyond vague input like, “This report was interesting; thanks for creating it.” Fusion center thinking should produce intelligence that generates one of two responses:

“This report was what we needed to decide on X.”

“This report was NOT what we needed to decide on X – here is what we need for next time.”

We suspect that when business leaders start responding with clear, definite feedback of that kind, fusion center analysts will be able to focus on fewer, but more-important, intelligence requests.

More Frameworks, and a Caution

We occasionally encounter the misconception in intelligence circles that all analytic frameworks must have the specificity and literal “framework” formatting of something like an Alternative Competing Hypotheses (ACH) matrix. For those unfamiliar with ACH matrices, they outline and give weighted values to different pieces of evidence to assess the validity or plausibility of multiple explanations of an event.

Although ACH matrices are excellent tools in the right scenarios, a typical ACH spreadsheet takes hours to complete, and the effort is not always justified. ACH is only one in a universe of useful frameworks at analysts’ disposal. There are many ways to organize data for maximum insight and utility. For example, in many situations a response-driven framework, in which analysts highlight the assessments most likely to help with short-term actions, can be just as useful as an ACH analysis or a detailed Diamond Model without requiring nearly as much work.

With that in mind, we recommend that analysts have experience with three types of analytical approaches that focus not on displaying data in boxes, but on showing the logical direction in which information flows. The shorthand for these approaches is the “-ductives,” referring to deductive, abductive, and inductive reasoning.


- Deductive analysis identifies *logical* conclusions based on available evidence. It focuses on the inescapable connections between causes and effects, or between nodes in a matrix. Deductive analysis is useful when analysts deal with datasets that could not have been falsified, and thus allow high-confidence assessments about an attack timeline.
 - Most likely use case: “How did X happen?”

- Abductive analysis identifies the best *possible* conclusion out of all the ones available. This framework is useful for questions like, “Who would want to target our content management system?” ACH is a specific implementation of this framework, but it applies when analysts whittle down the list of all explanations to the one that best fits the data and probabilities. Abductive analysis helps analysts answer questions about the explanation that, being most likely, demands the most attention. However, if asked, a good analyst should also be able to answer questions about other possible explanations.
 - Most likely use case: “*Why* did X happen?”
- Inductive analysis identifies what the *whole* looks like based on access to a *sample*. When a piece of malware demonstrates a new capability and leadership asks, “Should we be worried about this capability for malware in general?” the analyst will probably have to use inductive analysis to make generalizations about a trend based on the individual event.
 - Most likely use case: “*What* can we generalize about threats, based on the specific X that we can see?”

There are many cyber-specific frameworks for organizing and analyzing indicators of threat activity, including the Pyramid of Pain, the Diamond Model, and the Cyber Kill Chain. Analysts should become familiar with the models that are relevant for the domains important to their fusion center. However, they should not neglect widely applicable frameworks that stress logical direction and critical reasoning. And managers should certainly decide on the frameworks that will be most effective in their environment before, say, demanding that analysts know the entire MITRE ATT&CK framework so they can align reported threats with individual MITRE ATT&CK techniques.

Chapter III

Fusion Center Staffing and Data Sources



Humans: The First, Best Data Collectors and Analysts

In this chapter we discuss key points about creating and running fusion centers, which also apply to promoting fusion center thinking when you don't have a formal fusion center organization.

We start with a potentially controversial observation: in a successful intelligence shop, *who* is collecting and analyzing

is more important than *what* is collected and analyzed. While decisions about data sources, platforms, and documentation are very important, a good analyst is comfortable hunting for the best information even without the latest and greatest tools. The best threat intelligence producers working today would probably have been as good 10 years ago, even without the many resources available to today's teams.

What are the qualities of a successful intelligence analyst? Here are five characteristics that show up consistently.

Self-aware

Good analysts are excellent at escaping bias or tunnel vision because they are predisposed to assess their own behavior unemotionally on a regular basis. A good rule of thumb is that people who cannot discern their own deficiencies and work to improve will probably fail at sussing out the strengths and flaws in others' analysis.

Prolific

We would always rather work with *too much* output from an analyst rather than *too little*. An analyst's mind is constantly on the go, and great analysts typically produce more information than necessary because they do such thorough and wide-ranging research. You can trim a verbose report much more easily than you can insert missing insights into a sparse one.

Curious

We have encountered applicants for intelligence analyst jobs who claim an interest in a cybersecurity topic but can't provide follow-on insights. An example would be analysts who talk up their attraction to APT investigations and then draw a blank about *which* APTs they'd like to track. Truly curious analysts don't stop at mere interest; when their interest is aroused, they take the initiative and keep digging. Look for people who have no trouble talking at length about what intrigues them!

Detail oriented

One of the crucial distinctions between a human and an automated analyst is that a thoughtful human cares quite a lot about the details. Good analysts pursue information of increasing granularity so they can be as specific as possible in their assessments.

Flexible

Humans are imperfect. To improve, we need to regularly confirm the validity of our perspectives, methodologies, and data. Analysts should be happy to admit mistakes, look at assessments from different perspectives, and scrap pieces of analysis that are no longer useful. To detect a lack of this quality, look for *defensiveness*, which you can spot easily if you ask an analyst to explain more than two or three errors in their writing.

In Chapter VII we go into detail about the kind of biases and logical fallacies that intelligence analysts need to avoid at all costs. You should look for these weaknesses when interviewing job candidates and take them into consideration. However, you can also provide education to help new and existing team members recognize and correct biases and faulty logic.

Four Major Types of Data Sources

Our next topic is how to feed knowledge to domains for successful convergent analysis. In other words, what data should we be collecting, and how should we collect it?

The more experience you have when attempting to articulate connections between cyber and non-cyber domains across multiple intelligence areas, the more nuanced these questions will seem. The truth is that analysts must constantly refine the requirements for identifying, collecting, aggregating, filtering, and interpreting data for convergent analysis. There are no silver bullets, no complete, out-of-the-box toolkits, and no substitutes for hands-on familiarity. On the other hand, certain tools and tactics absolutely perform better than others, which is why we will lay out several best practices over the next several pages.

Once you have assembled a team of analysts with the qualities described above, you can start to confirm and expand the sources of data that this team will use to produce fusion center analysis.

There are four major types of data sources we will discuss:

1. **Internal** sources for **cyber** data
2. **External** sources for **cyber** data
3. **Internal** sources for **non-cyber** data
4. **External** sources for **non-cyber** data

Start with Internal Cyber Data

The highest-priority data for a cybersecurity intelligence team is internal, security-related data. No amount of cultural or economic context can bolster a security program that lacks a clear view of the devices and networks it is set up to protect.

Log files

The types of data that are necessary for an advanced cyber threat intelligence operation are well documented. At the top are the numerous log files that monitor activity in software and hardware across the enterprise. [NIST Special Publication \(SP\) 800-92](#), “Guide to Computer Security Log Management,” gives extensive recommendations for the kinds of logs that teams should have access to and be able to understand.

NIST prioritizes security tooling logs – such as for antivirus programs, IDS and IPS software, web proxies, authentication servers, routers, and firewalls. While these are all table stakes for security data collection, we recommend that analysts first gain experience with the “default good” behavior of user machines through logs of host operating systems. (Anyone protecting an Active Directory network should, without question, be smarter about Event Viewer than any individual EDR or WAF tool.) Finally, NIST outlines several application logs to track, such as client requests, server responses, usage information, and “significant operational actions,” which will vary widely depending on the application and how normal user behavior has been delineated.

SP 800-92 does note that collecting a large volume of log data is unhelpful if the organization does not have the personnel and tools to manage it effectively: “The fundamental problem of log management is balancing a limited amount of log management resources with a continuous supply of log data.”

The need for effective log management is probably the most important driver of spending on tools for the security team. Allan Liska, one of the world’s foremost experts on ransomware and a long-time contributor to the cybersecurity field, sums up the situation like this:

Over the last 10 years, the Security Operations Center (SOC) has become a dumping ground. If there are logs, they go to the SOC. In multiple companies, if you walk in and scan your access card, those logs go to the SOC, and the SOC doesn’t know what to do with them... Log data has at least quadrupled in the last 15 years. I don’t know of anyone who says they’re collecting fewer logs. Maybe you didn’t collect Windows logs in 2010 - now you have to collect them, so much so that Microsoft charges extra for collecting certain types of logs.

Of course, the reverse situation is also common – and worse. Many organizations fail to gather logs from applications, devices, or networks.

There seems to be a threshold of security maturity. Teams under the threshold scramble to make up for a *lack* of data, while teams above the threshold labor under *too much* data. The former condition stems from a failure to collect logs effectively; the latter, from a failure to manage them effectively.

When security leaders evaluate solutions and services that organize and aggregate threat data, we recommend they look for tools that are strong in two key areas:

- *Automation:* Software that directly lists or aggregates log files should make it easy to create filters with associated priority levels and alerting behavior that can then largely handle the volume of logs independently. You want tools that don’t force analysts to spend too much time clicking through alert boxes and emails at the beginning of their workday.

- *Investigation:* Software should make it easy to conduct investigations into logs and other data at increasing levels of granularity. Analysts should be able to quickly get back summary answers and to search by particular attributes using syntax simple enough that they don't have to dig through 100 pages of documentation.

Security tools

After the nuts-and-bolts data from log files, organizations should track operational and administrative information associated with security. Once again, NIST provides a helpful resource: [SP 800-137](#), "Information Security Continuous Monitoring for Federal Information Systems and Organizations." Despite the "Federal" in the title, this resource contains useful lists of detection and management software for all organizations, including products to manage data about vulnerabilities and patching, events, incidents, malware, assets, configurations, networks, and licenses.

Continue with External Cyber Data

We have been discussing cyber data that resides inside an organization's perimeter. There is also a huge range of valuable cyber data available outside of the organization, but collecting and assessing external data involves a new set of challenges.

ISACs

Midway between purely internal and purely external data is the kind of intelligence available from industry- and region-specific information-sharing groups. Many label themselves as an Information Sharing and Analysis Center (ISAC), such as the financial industry consortium FS-ISAC.

Open source data

Analysts can make use of open-source data shared on numerous platforms, including:

- Blogs
- Code repositories
- Internet forums
- Mainstream news outlets

- Experts and online services that run malware sandboxes
- Paste sites
- Security vendor websites and newsletters
- Social media

Dark web sources like criminal marketplaces and underground forums about cyber threats should also be monitored when possible.

As with log data, security teams should pursue consolidated tools to effectively automate and investigate external cyber data.

Proceed to Internal Non-cyber Data

In the “Fusion Center Thinking Starts at Home” section of Chapter I, we mentioned that a critical source of non-cyber data is the business itself.

The most capable adversaries will study an organization to determine what they can steal from it, what employee and customer information can be compromised and further abused, what intellectual property would be valuable for competitors and nation-states, what disruptions of business and technical processes would be most damaging, and what disclosures of confidential information would lead to the most embarrassment.

To “think like an adversary,” fusion center teams need to gather knowledge about many aspects of their organization, including:

- Accounting and financial systems
- Processes for research and product development, manufacturing, logistics, supply chain management, sales and marketing, and other key business activities
- Interactions with customers and employees
- Valuable intellectual property
- Locations of operating units and key suppliers
- Competitors
- Leaders
- Reputation
- Involvement in controversies

To obtain internal non-cyber data, fusion center teams (and cybersecurity teams in general) need to take the initiative to establish relationships with relevant business units and staff organizations like HR, legal, and finance.

Dive Into External Non-cyber Data

Challenging as it is to collect cybersecurity-specific data, it is even trickier to identify the data from non-cyber domains that you need to stay ahead of cyber threats. Fusion center analysts need to address questions such as:

- What non-cyber domains should receive priority?
- What types of data should we monitor to assist with anticipating geopolitically motivated threats?
- What types of financial and technological intelligence have a strong enough connection to cyberattacks to warrant regular reviews by our fusion center?

Preparing to use external non-cyber data

The first step is to review whether your organization even has the capability to ingest non-cyber intelligence as context for cybersecurity decisions. This requires asking questions such as:

- Do I have robust, comprehensive coverage of internal cybersecurity data (e.g. firewall, OS, and network logs)

Comment: If you cannot track persistence or lateral movement in your network after exploitation of a public-facing router, no amount of insight into Chinese national objectives will help your data confidentiality!

- Is my organization adept at translating *internal* cybersecurity data into effective tactical and strategic action? For example, can my organization quickly act on an indicator of compromise? Can we make effective decisions about longer-term security investments based on intelligence assessments created via internal data?

Comment: The larger the organization, the more likely it is that many sources of data are underutilized and that disparate departments (e.g., vulnerability management and IT deployment) are not effectively talking to each other. Having the right data is the first step, but making good use of it across collaborative departments involves several more steps.

- Is my organization adept at translating *external* cybersecurity and non-cyber data into effective tactical and strategic action? What external data does my organization rely on? How much of this data is used versus simply owned? (i.e., Do I have many security vendors, or even just a few vendors, that provide more data than I can effectively consume?) What metrics do we use to *justify* the data we receive externally?

Comment: Just as simply having access to internal logs is not sufficient for deriving value from them, simply paying for security data or tools is no guarantee that they are helping a security program.

Information from external non-cyber domains is most valuable for anticipating triggers for cyber threats. But your organization will need to have its security house in reasonable order before you can take full advantage of expertise in external non-cyber activities. A cybersecurity organization in constant crisis management mode is not going to have the attention span to be creative and innovative in its use of external intelligence.

Selecting analysts from non-cyber domains

Our first recommendation is to hire people with existing expertise in fields relevant to your organization, who already have extensive insight into the issues and the best sources of data in the critical domains.

However, if you don't have the resources to hire best-in-class analysts for every topic area, hire great generalists rather than "good enough" specialists. It is typically only very lucky

and well-funded businesses that can hire excellent analysts in every domain, and strong generalists will be much more valuable as security and business needs change.

As we mentioned earlier, the most important non-cyber domain for any organization is its own business. That domain includes its industry, competitors, corporate makeup, proprietary products and technology, differentiators, leadership, etc.

You want analysts with the right backgrounds and aptitudes to perform both outside-in and inside-out analysis:

- Working from the outside in, they should be able to determine what external threat actors are most likely to know about your business based on the widely available public data.
- Working from the inside out, they should be able to deduce which of the business's data, services, technologies, and other sensitive and proprietary assets would be most valuable to outside attackers, as well as exactly why and in what respect.

Most organizations have a reasonably good handle on how certain types of information such as financial data and personally identifiable information (PII) can be monetized by criminal actors. However, to be comprehensively secure, they need analysts who understand the motivations of competitors, governments, activists, and insiders and the information assets they target.

Prioritizing data sources

What can organizations do to identify the most valuable sources of non-cyber data?

One useful activity for outside-in analysis is to have analysts and red teams work together on penetration testing. This exercise stretches their thinking about the organization as a target for a variety of adversaries seeking not just PII and financial data, but also software code, product designs, business plans, lobbying and government relations activities, logistics and supply chain information, and other “crown jewels” that otherwise might not be considered by cybersecurity specialists.

In contrast, inside-out analysis may not require much external data, but does demand input from every part of the organization.

A quick-and-dirty way to educate security teams about the rest of the business is through thoughtfully constructed brown bag presentations. Multiple departments can explain what a “day in the life” looks like for individual contributors and managers.

A more thorough approach is to ask each department to conduct a “process audit” that maps all their activities. This audit will bring to the surface all types of confidential information and potentially vulnerable processes. It will force analysts to determine all the data and systems that need to be protected, and may suggest ways to improve cost-effectiveness, security, and compliance (e.g., by anonymizing data or consolidating it into fewer, more-secure repositories). We suspect that such audits can also be helpful for senior managers who want to understand how teams are spending their time.

Of course, specifics will vary by industry:

- Manufacturing organizations must understand what is proprietary about their factory operations, technical product documentation, and supply chain.
- Media firms need to track what in the public sphere is most likely to invite controversy or partisan opposition.
- Software publishers and cloud services providers must think about how their code is developed, where it is stored, what third parties rely on it, and how it can be updated securely.

Acquiring and managing external non-cyber data

Today, so much data is available that many organizations of all sizes choose to work with information aggregators and intelligence platforms. These collect information from a wide range of sources and give analysts tools to filter and search data to meet the needs of ongoing intelligence requirements and specific queries.

For example, there are excellent services that aggregate data and news for financial institutions and other highly regulated organizations.

We've had success using free platforms like Reddit and Feedly to track events within our domains. At first, these forums appear very noisy and dominated by irrelevant information and opinions. However, good analysts eventually learn what news is worth tracking and how to filter for it.

Threat intelligence platforms are also available. (*Shameless self-promotion:* Recorded Future, our company, offers a platform that categorizes and enriches both cyber and non-cyber intelligence from thousands of sources and provides tools to automate and optimize threat intelligence workflows.)

Over time, good analysts start to recognize both good *sources* of data (say, a particular journal) and good *sharers* of data. Certain social media users draw on an impressive network of sources to obtain and publish a lot of great information on a daily basis. One of our favorite sharers of cybersecurity reporting is Catalin Cimpanu (@campuscodi.risky.biz) on BlueSky. Analysts should keep an eye out for these accounts and note the sources they rely on.

A few more thoughts

Fusion center analysts should be able to support each other, so a certain degree of overlap in knowledge and responsibilities is desirable. As Maggie McDaniel, the vice president of Insikt Group at Recorded Future, said:

In terms of coverage and portfolios, you want a level of redundancy built into your [intelligence] program. You don't want everyone to be a broad generalist. You want people to have one primary area of responsibility and some secondary areas, so someone's primary could be someone else's secondary. So when you map out intelligence requirements to figure out coverage, you can map the subject areas, like a mosaic model of primary and secondary functions, and then you move backwards into the technical skillsets that are required to cover those. This then applies to different backgrounds and tenures in the field.

Artificial Intelligence: Complement, not Substitute

Should you be interviewing AI agents to fill the positions in your fusion center?

In theory, advancements in generative AI should mean that intelligence teams will have much less to do. In practice, although AI can automate parts of the intelligence process, it should be employed as a complement to the work of human analysts and not as a substitute.

One need only spend a few hours with a common generative AI platform like ChatGPT or Claude to recognize their strengths and weaknesses in the context of intelligence.

Generative AI is very good for:

- Reviewing large sets of structured data (e.g., tables) and unstructured information (e.g., transcripts of conversations) at high speed
- Summarizing information
- Aggregating data and assessing patterns across multiple sources

But AI often stumbles when attempting to perform some of the most critical intelligence tasks:

- Identifying the most important, unusual, or relevant pieces of information based on intelligence requirements
- Anticipating what data is needed beyond the scope of existing sources (which goes to both creativity and expertise)
- Responding intelligently to feedback about gaps and errors in existing methodologies
- Providing specific rather than generic assessments

This is not to say that AI cannot be trained to improve its responses within certain narrowly scoped areas of subject matter expertise. However, AI cannot, by its very nature, “think outside the box” of the material on which it has trained.

In contrast, even the most junior human analysts can quickly recognize when they need extra data and support, and can imagine ways of answering an intelligence question using information and methodologies from other domains.

Our recommendation is to use AI as a very fast first-draft writer of intelligence outputs. But AI should *not* be relied upon to create high-fidelity intelligence, nor to produce the final versions of analyses and reports.

Call us old school, but we believe that those who can't write clearly and logically won't be able to produce clear, logical writing just by prompting an AI agent. People who can't assess a situation and recommend actions on their own are unlikely to produce high-quality intelligence and insights by using AI as a crutch.

Over the past year, we have seen encouraging patterns in AI usage for data collection and summarization. For example, in May 2024, the CIA [reported](#) “home run” success with an open-source large language model (LLM) called Osiris, which the agency used to summarize huge volumes of information. In December 2024, Google released [Gemini Deep Research](#), a research tool for the company's Gemini Advanced 1.5 Pro platform for generative AI. Deep Research can formulate a research plan for a piece of analysis, collect data from what it considers relevant websites, and organize the resulting findings into a document with a summary, sections, and citations.

When we experimented with Deep Research, a few first rounds with the tool produced impressive results. We can imagine how the powerful data collection, aggregation, and summarization capabilities of Osiris and Deep Research could save us days when creating situation reports on novel and unfamiliar topics.

However, these tools suffer from issues that are either inherent to AI or are very difficult for it to overcome. For example:

- Current-generation AI models are bad at understanding the differences among sources, and would likely struggle to evaluate new sources based on criteria obvious to a human analyst (such as the motivation of an expert author writing a blog post versus a vendor employee publishing an article designed primarily to

drive traffic to the company's website).

- AI models such as Deep Research attempt to incorporate input from more documents than necessary because they do not know how to prioritize data from a few highly reliable sources.
- Even with specific prompting, AI tools struggle to isolate and highlight the elements of aggregated information that would be most useful and interesting to a human reader — elements that human analysts put in executive summaries and “bottom line up front” sections of reports.
- Much like a sixth-grade essayist, AI models are very good at listing facts about a topic, but often cannot explain why its findings are meaningful and actionable or address the “so what” questions that human analysts anticipate and answer in their analyses.
- AI tends to deliver assessments that are highly generic, based on the broadest and least-interesting patterns it can easily identify.
- Models like Deep Research have to be strictly trained to avoid sourcing from data generated by other AI tools, which has not been vetted and may be the result of “hallucinations” (a situation sometimes referred to as “AI feeding on its own slop”).

In our experience, generative AI is much better at finding *relevant* sources of information than at effectively evaluating the *validity* of these sources or the *actual content* they are providing. For example, when we queried a platform called Perplexity about “the best domains for a cybersecurity analyst to be comfortable with that are not within computer science,” it returned a lot of results that needed further vetting, including some that were simply wrong or not responsive to the original question. A Perplexity query was like a faster and more organized Google search, but by itself it was not a good way to assemble meaningful data and statistics.

To be fair, human analysts sometimes also make the mistakes discussed above. But fusion center managers can train their analysts to recognize and avoid these mistakes, while AI companies may not be motivated to do the same with their models.

Chapter IV

Leadership, Communication, and Technology



Words like “fusion” and “convergent” sound compelling. Who doesn’t want to create a #bettertogether story that improves security and reduces risk? However, it can be a challenge to manage the inherent organizational complexity of fusing two or more subject domains and their workflows. This chapter examines some of the organizational ingredients needed to execute a successful fusion center.

Leadership

A strong executive sponsor is critical to any collaborative effort. By strong, we mean a leader ⁷ with the vision, “juice” (influence), and people skills to drive change, obtain the necessary resources, and become a primary cheerleader.⁸

Vision

Enterprises often tilt toward organizational complexity. Typically, this is because years of growth (or stagnation) produce a jungle of titles and overlapping functions that create a bias for inertia.

An effective executive sponsor can shield a fusion center from internal politics and work with other senior leaders to help fusion center participants understand how their work contributes to a larger purpose. The vision should help motivate team members, align them with a mission, and answer the ever-important “why?” questions.

Collaboration

The goal of any fusion center is, of course, to combine operational functions, which improves time to action and reduces risk. Strong leaders create an atmosphere of collaboration, where success is the result of a unified team effort and credit is shared. Anyone can relabel a SOC as a “fusion center,” but to truly blend disparate functions, fusion center leadership must have broad buy-in from adjacent peers (e.g., managers of areas such as security architecture, engineering, and incident response).

Resources

Fusion centers also require resources. People and technology aren’t cheap. Great architects, engineers, and analysts are difficult to find and retain. Attracting them requires competitive compensation and a compelling fusion center vision and goals. When requesting resources, an effective executive sponsor will evangelize the fusion center’s ROI up front, helping justify long-term investment and protecting the budget when cuts are proposed.

7 <https://hbr.org/2023/12/8-essential-qualities-of-successful-leaders> (requires subscription)

8 <https://hbr.org/2008/06/why-the-celtics-wonlessons-fro> (requires subscription)

Cheerleading

Motivation may represent the most important aspect of fusion center leadership. Expressing gratitude for individuals and their contributions is fundamental to leadership. This is especially true in fusion centers, where the best teams work long hours because they love the camaraderie and atmosphere.

Communication

Effective communication skills are no longer optional in business.⁹ They're a fundamental requirement for success. Good communication is even more important when operational defenders are regularly working together across different domains. The cybersecurity domain relies on computer science terminology, geopolitical analysts have their own dialect, with a style that can veer toward the verbose and academic, and fraud experts maintain their own lexicon of acronyms and abbreviations. Fusion center members must master verbal and written communication skills that convey the insights from each realm while maintaining simplicity and clarity.

Additionally, in this post-pandemic era, intelligence professionals may have few opportunities to collaborate in the same physical space. Virtual fusion centers demand even better communication skills.

Hiring for Fusion Center Thinking

Understanding how a candidate thinks, processes information, and solves problems is helpful for determining cognitive abilities. There are plenty of resources that can help you design interview questions in that area. You can't rely only on lists of accomplishments and accolades in resumes.

Collaborators, not peacocks

It is equally important to assess the potential for cultural team fit. For example, in our experience, nothing will destroy team morale quicker than hiring a "peacock," a person who may be very talented but who demands constant personal recogni-

9 <https://hbr.org/2022/11/how-great-leaders-communicate> (requires subscription)

tion and affirmation. These individuals are anathema to high-performing fusion centers, which thrive on collaborative accomplishments and collective recognition.

To ensure a good team fit, ask candidates and their references questions that probe (subtly) for peacockery. For example, when queried about accomplishments in prior roles, does a candidate respond more often with references to “we” or “I”? There’s nothing wrong with “I.” However, if there are few “we” references in relation to large accomplishments, you should dig further with follow-up questions such as: “How many people worked on that project?” and “What was your role in that initiative?”

Make extensive use of reference checks

Reference checking is an essential element of hiring for fusion centers. AirBnB founder Brian Chesky articulates well the need for thorough reference checks.¹⁰ References provide less-filtered feedback and offer an opportunity to ask savvy questions and to uncover additional perspectives on a candidate. You might ask questions like, “Why is this candidate amazing?” and “What specifically did they achieve?” Questions like “Who is the best person you have ever worked with in this type of role, irrespective of this particular opportunity?” will tell you what characteristics the reference values.

Ask references for instances where a candidate demonstrated excellent communication and collaboration skills and fostered teamwork that produced successful outcomes. You might also ask about scenarios where a candidate failed to communicate well; these are potentially more revealing than effusive praise.

Also ask references to provide concrete examples (ideally including numbers or metrics) of situations where candidates demonstrated subject matter expertise, achieved positive outcomes, or improved methods for intelligence collection, analysis, or reporting. To help frame those questions correctly, consult with other managers who are deeply familiar with the topics, workflows, and expected deliverables for the open positions.

10 <https://youtu.be/xpp2hawRbAc>

Technology

Once an organization decides to invest in a fusion center, quick results are paramount to ensure continuing management support and ultimately success. For that reason, it is critical to tackle the topic of technology and tools as soon as a fusion center team is in place.

In this situation, time is of the essence and perfect is the enemy of the good.¹¹ Although it may be tempting to try, you are not going to quickly build from scratch a comprehensive, scalable solution for data wrangling and analysis, intelligence assessment and production, result sharing, etc. Either buy an off-the-shelf product, subscribe to an online intelligence platform, or, if you feel you must create your own solution, start with a “minimum viable product” and build on it over time. Many a good intention has died because CEO patience isn’t perpetually elastic.

It’s also easy to fall into the trap of excessively testing third-party options. In any software category there are likely to be 10 or more pre-built solutions, but you don’t want to get bogged down in the selection process. Your organization may already have a software solution that will immediately solve for 90% of a fusion team’s objectives.

For example, it might be tempting to test products like Clickhouse and Databricks for database storage and management. However, Confluence may suffice for short-term collaboration while you pursue a long-term solution.¹²

Although GenAI developments are making headlines weekly, the implications of incorporating AI into workflows are still only vaguely understood. Without question, innovation will remake the landscape for fusion center workflows. While JIRA tickets may suffice in the short term to support fusion processes, in the long term all security-based workflows will undoubtedly be based on technology stacks that include GenAI. Agents will train agents on workflows and private LLMs will likely act as the primary interface for collaboration with both machines and fellow fusion center humans.

11 Alternate maxim, from China: “Better a diamond with a flaw than a pebble without one.”

12 <https://clickhouse.com/>, <https://www.databricks.com/>, <https://www.atlassian.com/software/confluence>

The most important long-term topic for personal development and growth, particularly in analyst workflows, is GenAI. LLMs are advancing at a mind-boggling pace. Prompt engineering and agent training skills are quickly becoming table stakes in cyber functions. This new industrial revolution need not spell the end of fusion center jobs; instead, harnessing the best LLMs to automate workflows will free human brains for higher-order tasks. Analysts and engineers must invest time to remain current on the best LLMs for different tasks and the best tactics to optimize LLM results.

Chapter V

Working Examples – Fraud Fusion Centers

Today, fraud fusion centers are the most widely implemented form of fusion center in the private sector, so in this chapter we investigate some examples.

Gartner estimates that 5% of large enterprises currently maintain cyber-fraud fusion centers.¹³ That figure is expected to jump to 20% by 2028. The reason for that growth is not hard to grasp. In January 2024, Nasdaq released a remarkable report¹⁴ with a finding that “Losses from fraud scams and bank fraud schemes accounted for nearly \$500 billion globally in 2023.” That number is a significant enticement for threat actors to continue their successful tactics and invest in innovative new techniques for perpetrating fraud.

A successful fraud fusion center removes data silos and works across teams that are traditionally separate, such as cyber threat Intelligence and fraud prevention. Fraud fusion centers improve communication and collaboration among teams through organizational (and sometimes physical) proximity, leading to better threat actor campaign analysis and reduced losses.

Most financial services companies create annual plans articulating objectives for fraud losses. Everyone is happy when they meet those goals. However, merely meeting expectations potentially leaves tens of millions of dollars or more on the table. The objective of fraud fusion centers is to apply

13 <https://cxotoday.com/specials/gartner-why-cyber-fraud-fusion-is-the-future-of-online-fraud-detection/>

14 <https://ir.nasdaq.com/news-releases/news-release-details/nasdaq-releases-first-global-financial-crime-report-measuring>

knowledge of multiple domains to reduce fraud to as close to zero as possible (though annual variable compensation targets are more realistic). That's an ambitious goal, but its simplicity helps unify everyone's efforts.¹⁵

A thriving fraud fusion center requires a leader with experience in both cyber defense operations and fraud prevention, as well as the authority to remove traditional barriers to collaboration across disciplines. These barriers include organizational empire building, information gatekeeping, and vendor firewalling (denying adjacent departments or groups access to information from a third-party vendor). Since fraud prevention and threat intelligence typically have different organizational reporting paths, a transformative leader who can bridge this gap is an indispensable factor for success.

In most organizations, CTI analysts focus on identifying malicious fraud communities and their TTPs. Fraud analysts are usually concerned with perfecting fraud analytics. In traditional organizations, the CTI and fraud teams sit in separate departments and have few incentives to share information or jointly devise improvements. Nobody is tasked with optimizing information workflows for maximum fraud reduction.

When implemented correctly, fraud fusion centers remove counterproductive incentives and enable a unified team to make the organization more effective.

Fraud Fusion Centers in Action

In the following examples we'll examine how cyber threat intelligence can enrich fraud investigations and help shift fraud reduction activities from reactive to proactive. We'll also discuss opportunities for innovation, collaboration, and communication, and for providing additional deliverables.

Fighting payment card fraud

A common objective of fraud fusion centers is reducing payment card fraud, including both the "card present" (CP) and "card not present" (CNP) varieties.

15 For more insights about fraud fusion centers, see <https://go.recordedfuture.com/hubfs/white-papers/castle-dilemma-cti-fraud-fusion.pdf?hsLang=en>

The first step of CTI teams is typically to identify criminal forums, marketplaces, and other online communities where stolen payment card details are advertised and sold and collect copies of those details. Unfortunately, collecting data from Telegram channels and dark web marketplaces is not as simple as sending a query. Criminal communities often screen participants or require tailored “bots” capable of remaining in the community forum and collecting unstructured text.

Cyber threat intelligence teams can also store and query payment card numbers and associated metadata such as expiration date, card verification value (CVV), and cardholder address.

Fraud team activities may include:

- Identifying at-risk cards and monitoring the associated portfolio exposure
- Tracking and reducing fraud mitigation turnaround times
- Reducing criminal interest in the payment cards of a specific institution (its payment card portfolio)
- Finding ways to decrease false-positive rates in risk-scoring solutions
- Pinpointing compromised merchants through common point of purchase (CPP) analysis (that is, by detecting patterns of fraudulent activity associated with specific merchants)

Commonly used fraud controls include:

- Enforcing know your customer (KYC) rules
- Monitoring fraud precursor signals (for example, using test transaction monitoring, which tracks nominal charges on stolen payment cards fraudsters use to ensure the cards’ viability for large fraudulent charges)
- Investigating alerts from compromised account management system (CAMS)
- Implementing CPP analysis.

However, all of these activities are reactive. Fusion center thinking changes the paradigm so that intelligence drives fraud prevention, mitigates the efficacy of fraudster tactics, and reduces reliance on fraud controls as a first line of defense.

One example of how cyber and fraud teams can contribute to each other's successes is in the area of stolen payment cards:

- Fraud analysts can help CTI analysts better understand bank identification number (BIN) targeting by threat actors and successful monetization trends. This perspective might help hone CTI data collection. For example, if fraud is moving away from the online channel to phone purchasing, CTI analysts might be able to adjust their data collection efforts to identify new underground services for recruiting phone operators who speak fluent English. Or, if a particular BIN is experiencing higher fraud losses than normal, the CTI team can work on identifying the threat actors responsible and documenting their TTPs.
- Then the CTI analysts' findings could be funneled back to the fraud team for fraud control tuning, allowing the team to increase scrutiny of the transactions for the most at-risk merchants and industry verticals.

When it comes to fighting payment card fraud, fusion center thinking can become a flywheel that accelerates learning and action as more insights are exchanged. We say "insights" because, while the underlying data should already have been available to all fusion center participants, the insights might never have been shared.

PII theft and ATOs

Some organizations, such as retailers, experience heavy fraud losses related to payment cards even though they don't issue cards themselves. These losses relate to the theft of PII, including physical and email addresses, birth dates, Social Security and phone numbers, and maiden names. Threat actors use this PII to tailor social engineering campaigns that compromise customer accounts in what are called account takeover (ATO) attacks.

Banks are also victims. They may be able to quickly identify and reissue stolen payment cards, but the associated stolen PII may lead to fraud losses through call center social engineering and online channel security bypass.

In another life, working with banks, we engaged fraudsters online to identify previously dormant cash-out accounts. At the time, even if a victim wasn't participating in online account access, fraudsters would enroll them using stolen PII and create the necessary authentication details (username, password, etc.). Out of convenience, all compromised accounts contained the same password.

Account takeover mechanisms include MFA prompt fatigue (the victim is tired of selecting “no” and seeing a dialogue prompt reappear), credential stuffing / brute forcing¹⁶, social engineering, SIM swapping¹⁷, password resets, etc. Naturally, GenAI is reducing the time to prepare and deliver attacks while increasing the quality of social engineering campaigns.¹⁸ In addition, ATO and associated cashout techniques are constantly changing.

To combat PII theft and prevent account takeovers, fraud professionals can collaborate with CTI peers to identify which threat actors to monitor and engage with. Engagement generally involves using an online moniker or legend that creates conversational opportunities within bad actor communities. Communication with actors is valuable because they often discuss and share tools and tactics. Given that there is an ocean of bad actors online at any given time, it is important to make good decisions about which to engage with. Fraud professionals can describe in detail how stolen PII is being captured and monetized, which helps CTI professionals narrow the type of actor communities for focused engagement.

Large databases of PII have historically been valuable targets for theft. Today, even after multiple arrests, infostealers (a malware genus focused on extracting victim PII from web browsers) have proven devastatingly effective – to the tune of billions of stolen victim credentials, cookie values, and more. This treasure trove of valid credentials enables fraudsters to quickly execute ATOs for social engineering attacks and to direct money movement within financial services platforms.

16 <https://therecord.media/hackers-breached-california-state-welfare>

17 <https://therecord.media/sec-twitter-account-hack-arrest-alabama>

18 <https://www.recordedfuture.com/research/qrcode-and-ai-generated-phishing-proliferate>

CTI analysts can engage the actors who own and operate info stealers to purchase compromised credentials. They can then feed those credentials back to the fraud team for proactive remediation before ATOs occur.

However, to monitor the info stealer ecosystem, it is essential to speak and type Russian. CTI teams should have a Russian-speaking analyst available to facilitate these long-term engagements. Separately, the fraud team can inform the fusion group about how and where stolen credentials are successfully being used to access client accounts.

PII is also at risk from malicious insiders.¹⁹ The dark web offers marketplaces where rogue employees and third-party contractors can sell customer PII. A CTI presence (both human and machine) in these marketplaces can enable an organization to locate stolen PII and help the fraud team determine if insiders are responsible for the theft.

Opportunities for collaboration don't end with fraud. For example, cybercriminals sometimes use ATO and cash-out services (compromised bank accounts) to facilitate money laundering. CTI analysts who gather intelligence on those activities from dark web and other sources can collaborate with anti-money laundering (AML) groups to identify and disrupt money laundering schemes – including some linked to terrorism groups.²⁰

Magecart compromised merchants

Magecart attacks use specific malicious scripts that, once loaded into vulnerable e-commerce platforms, skim payment card details during customer transactions. Magecart victims may have to consider more than just replacing their payment cards in [certain shopping contexts](#).

Proactively obtaining Magecart intelligence is important for e-commerce retailers and the financial services companies that advise and alert them, particularly because Magecart TTPs evolve quickly. For example, Recorded Future's Payment Fraud Intelligence (PFI) team have observed changes in the operation of Magecart attacks, such as:

19 <https://www.bloomberg.com/news/articles/2024-12-30/bank-insiders-are-leaking-data-on-client-accounts-as-scams-surge>

20 <https://intelligence2risk.substack.com/p/fraud-funding-terrorism>

- Using Google Tag Manager
- Exfiltrating payment card data to Telegram
- Hiding malicious scripts within scalable vector graphics (SVG) elements and using stenography
- Using compromised websites as proxies to distribute malicious payloads

CTI teams can alert fraud professionals to fast-changing Magecart TTPs and help identify victimized merchants before the monetization clock starts. That information can allow fraud teams to block or place fraud flags on attempted payment card transactions at compromised merchants. It also enables the teams to track monetization paths and timelines, for example, by observing whether a fraudster attempts to commit fraud via phone or online, at which retailer(s), and how soon after cards are compromised.

Scam merchants

As the name implies, scam merchants are ecommerce websites created for the express purpose of advertising spurious goods and services – which are never delivered. These websites look and feel like those of legitimate retailers, but the fraudsters who create them gather credit cards from unsuspecting victims and shut the sites down after a few weeks. Scam merchant schemes are easy to implement because the barriers to entry for obtaining a merchant account to process payment cards online are practically non-existent.²¹ Also, banks and merchant processors are not strongly motivated to perform due diligence on new merchant accounts, because more merchant accounts and more transactions lead to more fees.

CTI analysts and fraud teams can collaborate to create and implement comprehensive strategies to uncover scam merchants. Strategies to thwart them combine open source research, collaboration with industry partners, and advanced data analysis. Organizations can identify suspicious websites by monitoring social media advertisements, reviewing domain registration details, and tracking patterns in merchant account behavior.

21 <https://intelligence2risk.substack.com/p/fraud-funding-terrorism>

Further validation of potential scam websites can be achieved by correlating domain data with payment card data exposed on the dark web, identifying common patterns across fraudulent sites, and leveraging consumer reports of suspicious activity. By integrating these insights, fraud fusion center teams can avoid emerging threats and effectively mitigate payment fraud risks.

Android banking trojans



Mobile malware is particularly pernicious in retail banking outside of North America. Recorded Future observes thousands of malicious files with anti-virus signature names like Hydra, Hook, and Sova, and more emerge every quarter. For example, [Thailand has been alerting](#) citizens to the dangers of installing unverified Android applications, which could drain their bank accounts. While mobile malware generally contains [many features](#) (e.g., accessing location services, camera snooping, reading/sending SMS, etc.) that achieve device takeover (DTO), the most fraud losses are caused by [man-in-the-browser \(MitB\)](#) capabilities created during online banking sessions. This is because traditional [multi-factor authentication \(MFA\)](#) mechanisms become less effective when all communications channels can be intercepted and manipulated, leading to account takeover.

CTI teams can create alerts and dissect new mobile malware TTPs to help financial services fraud teams and government investigators focus on key threat signals without having to grope through ever-expanding clouds of “digital exhaust” (the totality of meta-data such as IP addresses, cookie values, and time zone settings that digital devices produce as they traverse the internet).

Additionally, real-time notification of new mobile malware samples with associated metadata improves bank prevention/detection systems via atomic indicators (e.g., user-agent strings and IP addresses) and internal log correlation.

Stolen checks

Physical (paper) checks are a relic of the banking system employed primarily in North America. Because these documents contain routing and account numbers, they are still heavily used in [fraud](#), particularly against older people who [prefer checks](#) over newer [fintech options](#).

CTI analysts with access to criminal online communities can use [optical character recognition \(OCR\)](#) to review checks found there and alert fraud teams so they can forestall the fraudulent use of those checks, even before the check owners report an issue.

Chapter VI

Risk Assessment with Fusion Center Thinking



Risk Assessment Versus Forecasting

There are subtle but extremely important differences between *forecasting*, as practiced by meteorologists, media pundits, and futurists, and *risk assessment*, as performed by intelligence and business risk assessment groups. It is critical to understand these differences when setting expectations and creating goals for fusion centers.

Don't cyber intelligence groups forecast?

When a company shows up in the news as the victim of a massive data breach or espionage intrusion, it's not unusual for outside observers to criticize it for a lack of foresight: "They should have seen it coming." And indeed, much of the language around the benefits of cyber intelligence – staying "ahead of the adversary," anticipating "emerging threats," being "proactive rather than reactive" – implies that analysts should predict the future. If intelligence can't identify what's coming, what is it good for?

Of course, "to forecast" can have several meanings, but the most common one in regular speech is, to use the first definition listed by Dictionary.com: "to predict (a future condition or occurrence); calculate in advance: *to forecast a heavy snowfall; to forecast lower interest rates.*"

Hundreds of people make a living forecasting likely outcomes in the future. Most face few repercussions if their predictions turn out to be wrong. In contrast, businesses can ill afford to subsidize intelligence shops that make consistently inaccurate predictions. They also can't afford forecasts that cover the most likely future conditions or occurrences and nothing else.

Why? Because businesses (and governments and militaries) need to *identify and prioritize multiple risks*, for two reasons:

- In most situations, there is a high probability that several negative outcomes will occur instead of or in addition to the most likely one.
- Less-likely occurrences may have far greater impact on the organization than the most likely one.

In practice, the goal of most intelligence efforts is *reducing the totality of downside risks and increasing the sum of upside risks*. Insight into the future is meant to avoid multiple bad outcomes and lock in many good ones.

Too often, organizations rely on cyber threat intelligence to provide a single, well-hedged answer to questions about what will happen to a company or industry or country in the future.

For example, this is a fatal flaw in the otherwise excellent methodology proposed by the famous book *Superforecasting*. The desire to prioritize only the most likely outcome is not an inherently wrong impulse, but it is inherently incomplete. Organizations need to be preparing for several likely outcomes, as well as other events that may not be as likely but *might cause severe losses*.

Two scenarios illustrate this point. Someone who keeps their medicine cabinet well stocked with cold and allergy remedies, but smokes two packs of cigarettes a day, is well prepared for common negative occurrences but also exposed to a catastrophic risk (fatal lung cancer). Inversely, a person who builds a disaster-proof shelter but does not own an umbrella will avoid losing everything in a single day, but will get wet every time it rains. Both are imbalanced approaches to the future. To prepare for the average threat and the worst-case threat, fusion center analysts need to assess threats based on both the likelihood of their occurrence *and* their severity if they occur.

We expressed this point in several ways because we've seen how often it's missed or ignored by producers and consumers of threat intelligence. An organization that only prepares for its worst possible threat — say, nation-state-level malware targeting industrial control systems — might mitigate that threat by spending tens of millions of dollars, but face unacceptable accumulated costs from garden-variety credential leaks and cloud storage exposure. An organization with the most up-to-date firewalls and email filters might stave off 99% of threat activity, but suffer immense losses if an employee with access to business-critical data is compromised because of a lack of security awareness training and inadequate monitoring.

How intelligence producers and consumers respond to the combination of most common and most severe outcomes should be driven by an understanding of acceptable losses. As our co-author Levi Gundert wrote in his book *The Risk Business*²²:

22 Levi Gundert, "The Risk Business: What CISOs Need to Know About Risk-Based Cybersecurity," available at: <https://go.recordedfuture.com/the-risk-business-second-edition>

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, where 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?

These two terms – most common, most severe – are the backbone of appropriate fusion center thinking. “Forecasting” is one way of describing predictions of the future, but it can be misleading if the audience thinks of that term in the sense of “forecasting a heavy snowfall.” The more appropriate descriptor is risk assessment.

Fusion Center Thinking and Risk Assessment

Preparing for the most likely and the most severe outcomes requires analysis and prioritization. That’s exactly where risk assessment comes in.

Risk assessment is about:

- Systematically identifying *all* the most *common* or *likely* threats, plus the most *severe* threats, to the organization
- Estimating their probabilities
- Estimating their impacts
- Helping to find the most cost-effective ways to reduce overall risk

Cyber risk assessment is the mapping out of the most likely and most damaging futures in which a company suffers **downside risks** such as brand impairment, competitive disadvantage, compliance failures, financial fraud, and operational downtime because of cyber-related activities.

In our experience, fusion center thinking plays a critical part in cyber risk assessment because it leads to an accurate understanding of context-driven *triggers* for such threats. Knowing the most common economic, technical, and political drivers for cyber threats gives analysts an edge in determining the probability of different attacks and estimating the severity of their impact on individual entities and on financial markets, governments, and other significant collective entities. That knowledge enables organizations to make optimal decisions about which threats they should prepare for and which measures they should use for detection, prevention, and mitigation.

What to Ask from Fusion Center Assessments

You should have appropriate expectations of what it means for fusion center intelligence to be “forward thinking.” We cannot emphasize enough that, in 99% of cases, proper risk assessments are not forecasts that allow organizations to anticipate all the details of every attack, such as the timelines of threats affecting them. Outside of very particular examples, organizations should not hold their intelligence teams to a standard of predicting discrete events at specific times.

In line with good risk management, we propose that organizations use forward-thinking threat intelligence to answer four straightforward questions. These answers are not about *telling* the future (i.e., forecasting what will or won’t happen at a binary level) but about *preparing* for the most likely and most severe futures.

- Question 1: What are the areas right now where a threat actor targeting you *could* get in without mitigation or detection?
- Question 2: What are the areas right now where a threat actor who got through would cause the *worst effects*?
- Question 3: What are the most *common types* of current threat activity targeting your organization?
- Question 4: Of known threat groups, which are most likely to target your organization based on their *historical choice of victims*?

Answering question 1 requires *auditing* so you can understand the entirety of your organization's infrastructure and vulnerabilities. Your answers may influence assessments of future risks, but they are about determining the situation in the *present*.

Question 2 is also about auditing, with an extra factor: answering it requires the ability to *prioritize* elements of your organization's assets (data, networks, operations). This goes back to the discussion in Chapter III about making sure your security intelligence producers have enough knowledge about the larger business so they know where the greatest damage or disruption could occur.

Question 3 is important to ask so that analysts and intelligence consumers avoid the temptation to focus on forms of threat activity that are most novel (and thus likely to end up in the news and in security vendor reporting) but less likely to be encountered in reality.

(A quick aside: if your security intelligence producers are saying that the most likely threat to your organization in the next six months is a ransomware attack, they should go back to the drawing board. The most *likely* threat — even if it is not the most *severe* one — is almost certainly spam emails, automated vulnerability scanning by a threat actor, or access to a misconfigured database.)

Question 4, finally, calls for the most sophisticated and considered analysis. It requires the greatest expertise to answer correctly, so it should be given to analysts who have the most experience with threat actor motivations and background, and who are best able to avoid the biases and logical fallacies described in the next chapter. This question requires a *deep understanding* of the connection between threat actors' targeting patterns and your organization's public presence and assets.

These questions should be examined in a two-step process:

- First in the light of internal cybersecurity measures and knowledge of the organization
- Then in terms of the triggers in external non-cyber domains

Also, each of the four questions should be addressed in terms of specific non-cyber triggers.

For Question 1, the non-cyber domains most likely to be useful are associated with *technology and business operations*. Determining how a threat actor might be able to get in without detection often depends on knowledge of an organization's IT infrastructure. These changes can align with developments in corporate policies, third party relationships, individual employees' use of social media, etc.

The non-cyber domains most likely to affect answers to Questions 2 and 4 are associated with *geopolitics, economics, sociology*, and an organization's *competitive position*. Typically, threat actors seek to achieve their goals by discovering and exfiltrating valuable information, or by disrupting operations for reasons of ideology, or by eliminating competitive advantages. Analysts with the right domain knowledge can anticipate which of these value propositions is most relevant to specific threat actors, based on expert understanding of the different motivations of adversaries such as nation-states, competitors, criminal gangs, and ideologically driven individuals and groups.

Finally, the non-cyber domains that are most likely to help with Question 3 are associated with *economics, public policy, software development*, and *IT market share*. The prevalence of a threat type typically correlates with factors such as availability of targets, ease of execution, level of profitability, and ability to evade punishment.

Fusion center analysts with knowledge of multiple domains are in the best position to weigh these factors. For example, a certain intrusion tactic might be hard to pull off and invite immediate governmental scrutiny (negative factors for threat actors), but these drawbacks might be more than offset by access to very sensitive data from thousands of target companies (a huge incentive).

Fusion center analysts may also be in the best position to judge the effect of evolving non-cyber factors on changes in criminal behaviors. Examples include increased activity by some law enforcement agencies (such as American and European agencies shutting down spam networks); decreases in activity by other law enforcement agencies (such as Russia's growing tolerance of ransomware operators); growing unemployment (making it easier for criminals to recruit money

mules); increased availability of anonymization tools (making criminal communications harder to track); and democratization of advanced hacking tools (which opens opportunities for less-advanced threat actors).

We hope this discussion communicates the benefits of fusion center thinking in assessing threats to an organization. Anticipating cyber threats by identifying their *triggers* early, rather than waiting to respond reactively to their *occurrences*, saves organizations time and money and improves their security posture.

Risk Assessment Frameworks for Fusion Center Thinking

Framework #1: The CMMC Cybersecurity Maturity Model

One of the takeaways from our research into security architecture is that guidance from authorities is very useful and also rather punishing to read. For example, the U.S. Department of Defense's breakdown of requirements for the [Cybersecurity Maturity Model Certification \(CMMC\) Program](#) is an excellent framework for helping organizations fully audit their threats, vulnerabilities, and defenses, but nobody enjoys studying it.

There are two guidelines from NIST that do a good job explaining what high-quality predictive analysis in a fusion center looks like. The first is [NIST Special Publication \(SP\) 800-171](#), ponderously titled: "Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations" (translation: protecting confidential data in the private sector). Not only does NIST SP 800-171 describe the highest level of maturity to aim for within the CMMC program, it also outlines what an advanced risk assessment program should look like:

Periodically assess the risk to organizational operations (including mission, functions, image, or reputation), organizational assets, and individuals, resulting from the operation of organizational systems and the associated processing, storage, or transmission of CUI [Controlled Unclassified Information].

Clearly defined system boundaries are a prerequisite for effective risk assessments

[emphasis added]. Such risk assessments consider threats, vulnerabilities, likelihood, and impact to organizational operations, organizational assets, and individuals based on the operation and use of organizational systems. Risk assessments also consider risk from external parties (e.g., service providers, contractors operating systems on behalf of the organization, individuals accessing organizational systems, outsourcing entities). Risk assessments, either formal or informal, can be conducted at the organization level, the mission or business process level, or the system level, and at any phase in the system development life cycle. [SP 800-30] provides guidance on conducting risk assessments.

Pay attention to that bolded line about “clearly defined system boundaries.” To best understand the active and potential threats to an organization, and the ways that those threats could go unseen, fusion center analysts need to be thoughtful about how the organization is structured and what fits inside it. If your organization is thinking about “system boundaries” purely in terms of the nuts and bolts of IT infrastructure, then it is likely ignoring all of the ways in which payroll policy, or marketing events, or executives’ social media are part of a larger threat landscape. Determining these system boundaries should be an all-company affair, not merely a siloed project for already overburdened security teams.

Recall that in an interview mentioned earlier, a lead IT analyst for a manufacturing company talked about how the department responsible for trademarks was not just a critical area of the business, but also very important for determining the legitimacy of *newly* registered domains. SP 800-171’s advice urging readers to think about constantly widening the aperture regarding risk and impact, first at the level of internal assets and systems, and then at the level of external entities like suppliers, is excellent. We would add that fusion center analysts should think about the attributes of the wider business, its industry, and regions in which it operates to best understand what threats could be incoming.

Readers may have spotted a quick reference at the end of the quote above to another NIST publication, [SP 800-30](#), “Guide for Conducting Risk Assessments.” This publication goes even further in talking about the need for a wider group of stakeholders in forecasting threats: “Since cost, timeliness, and ease of use are a few of the many important factors in the application of risk assessments, organizations should attempt to reduce the level of effort for risk assessments by sharing risk-related information, whenever possible.”

SP 800-30 comes intriguingly close to calling for the kind of fusion center thinking we have been describing when it states that, in evaluating vulnerabilities:

Organizations [should] also consider predisposing conditions. A predisposing condition is a condition that exists within an organization, a mission or business process, enterprise architecture, information system, or environment of operation, which affects (i.e., increases or decreases) the likelihood that threat events, once initiated, result in adverse impacts to organizational operations and assets, individuals, other organizations, or the Nation.

SP 800-30’s example of a predisposing condition is putting an office in a “flood-prone region.” Other scenarios that could easily fit the definition of a predisposing condition for vulnerabilities in a business’s assets include:

- An office location in a country that has just declared war on its neighbor
- Reliance on certain forms of physical identification to create customer accounts
- Public support for controversial political figures or issues
- Creation of a new domain for a one-time company event
- Rapid acceptance of a new technology or platform for handling customer data
- A sudden spike in prices for consumer staples
- Sudden awareness of a brand due to unplanned social media virality

SP 800-30 correctly prioritizes attention to known threat actors. It recommends understanding the “capability, intent, and targeting” of threat sources when conducting a risk assessment. However, it doesn’t provide a framework for thinking about the most likely *triggers* for threat activity by known or unknown threat groups. Fusion center thinking can take larger environmental analysis into account when evaluating threats so that it doesn’t focus narrowly on intrusion sets that have been publicly reported.

Framework #2: RAND SWARM Analysis

In 2021, the RAND Corporation published a [white paper](#) regarding their “Scalable Warning and Resilience Model” (SWARM). SWARM is a framework designed to help organizations anticipate and address “the nation-state-sponsored threat class.” It involves four main steps: identify relevant cyber adversaries, focus on all-source intelligence collection, apply a threat model, and adversary emulation. Of these, identifying relevant cyber adversaries was the step most concerned with predicting future threats from past precedent.

RAND formulated a new predictive framework to make up for obvious misses on the part of intelligence functions in the past:

Despite increased focus on developing more-sophisticated cybersecurity tools and techniques for defending organizations against cyber threats—including using cyber threat modeling, information sharing, and threat-hunting—the cyber defense community is still largely reacting to, rather than predicting or anticipating, these threats.

By applying the SWARM methodology to themselves, and in particular by identifying the North Korean group Kimsuky as a “relevant cyber adversary,” RAND claimed to have successfully predicted a phishing email from Kimsuky within an expected 14-day window of news following talks between the United States and South Korea regarding sanctions on North Korea. Pinpointing an event so accurately is no easy feat, and RAND’s accomplishment here should go on record as among the clearest and best-defended justifications of heightening security measures in response to a geopolitical event.

The SWARM paper is very useful for scenarios in which organizations already have a good idea of which *specific and known* threat actors are more likely to target them. In the section on identifying relevant cyber adversaries, the paper states that organizations should use priority intelligence requirements (PIRs) to “facilitate the identification of known (those that have breached the networks of the organization in the past) as well as unknown (those that are familiar to the cybersecurity community but have not breached the networks of the organization in the past) threats to an information environment.”

Although SWARM was designed to help analyze threats from nation-state-sponsored groups, organizations can use it to evaluate other types of threat actors, including hackers, ransomware gangs, and initial access brokers.

On the other hand, SWARM is less useful for analyzing how non-cyber contextual factors could drive threat activity against high-value assets and networks, because it doesn’t take those factors explicitly into account as part of its framework for prioritizing threats. SWARM does not provide commentary on what other factors could inspire individuals or governments to engage in new cyber threat activity, beyond what has already been reported. In other words, SWARM’s analytical framework was not designed to consider an environmental analysis of the threat landscape.

Framework #3: The Threat Category Risk (TCR) Framework

To return to *The Risk Business*, in 2020 our co-author Levi Gundert proposed a framework for cyber threat risks incorporating many of the best practices that this book has already mentioned, such as appropriate attention to likely and severe outcomes, ability to defend assessments at a quantitative level, and consideration of external factors’ influence on threat categories. It is called the Threat Category Risk (TCR) framework because it evaluates how general *categories of impacts* associated with cyber threat activity (according to the traditional Confidentiality-Integrity-Availability model) arise from general *categories of threats* (such as social engineering or vulnerability exploitation).

For each threat category in the TCR framework, a team can estimate:

- 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 bounds of damage if a loss of confidentiality or integrity occurs
- The upper and lower bounds of the duration (in hours) and the upper and lower bound of the cost per hour if a loss of availability occurs

The value of the TCR framework aligns one-to-one with a business’s ability to articulate the damages associated with threat categories. In other words, this framework certainly benefits from the threat assessments generated by threat intelligence teams, but it only offers complete insights when those assessments are aligned with relevant costs. An organization that knows credential exposure is bad, but cannot estimate the cost of addressing exposed credentials, will find little value in the TCR framework.

Here we repeat a point made in Chapter I: fusion center thinking starts at home. Proper risk assessments must consider the unique resources and assets of the particular organization. A threat that could drive Company A out of business might only amount to a few days of cleanup for Company B. As implied by its name, fusion center thinking fuses data about an external threat landscape with data about an internal operational landscape. To maximize results, fusion centers need both kinds of data and the means to analyze them together.

A Final Example

Many forms of cyber threat intelligence are likely to feature assessments that seem to contradict some of our suggestions above. After all, doesn’t the following statement sound like a forecast of threat actor activity in a particular timeframe — a bit too close to fortune telling?

Within the next six months, we anticipate that threat groups sponsored by the Iranian government will be the most likely of nation-state groups to target our business as a result of our planned research and development of autonomous weapons.

In fact, this statement avoids many of the wrong ways to create assessments of this type (see Chapter VII):

- The assessment is not forecasting that a specific thing will happen at a specific time, but rather suggesting that within a given timeframe a certain occurrence is the most likely one. Good risk assessments use timeframes to contextualize assessments rather than to guarantee outcomes.
- The assessment prioritizes a threat *within a category*, not threats in general. It does NOT say that Iranian groups are the most likely *of all threat actors* to target the business. Rather, it states that, *if any nation-states* were to target the company, Iran would be the most likely for specific reasons.

This assessment is written appropriately because it takes contributing factors into account.

On the other hand, here is a bad risk assessment that looks like an unexamined forecast:

We anticipate that Iran will be very likely to target us in the near future due to our proprietary technology.

In this case, details necessary for decision making are missing. Fusion center thinking would produce a much more useful assessment.

Chapter VII

Critical Thinking within a Logical Framework

Intelligence is hard. It's certainly difficult to produce. The starting analyst realizes very quickly that there are layers and contingencies and exceptions to everything they're tasked to monitor. Fusion center intelligence adds another dimension of complexity by requiring the analyst to cross-reference different types of information from diverse sources. When examining cyber threats through the lenses of non-cyber domains, an analyst might pursue a tough research question for several days before realizing that, due to an inappropriate framework or misinterpreted data, like Sisyphus, they must start over at the beginning.

However, intelligence is also difficult to *consume*. Business leaders have to hope that the people delivering it understand the priorities of the business and know how to communicate at a strategic level, not simply sending up a list of technical observations. And leaders have to be concerned that intelligence producers might be watering down assessments because they fear to come forward with harsh but necessary recommendations. As the Harvard Business Review put it in 2004, "The CEO is often the most isolated and protected employee in the organization. No one gives him unfiltered information."²³

A final, critical challenge for intelligence groups is preventing assessments from being distorted by logical fallacies and unconscious cognitive biases. In fact, we're confident that the

23 <https://hbr.org/2004/02/worse-than-enemies-the-ceos-destructive-confidant> (requires subscription)

most important asset for effective intelligence production is not technical acumen or crisp writing, but rather the ability to think critically and objectively within a logical framework about the data being reviewed.

Now, we are not here to cover the same ground as a critical reasoning textbook. Plenty of excellent resources for understanding and avoiding logical fallacies and cognitive biases are available online and in bookstores. But we feel it is essential to give intelligence producers and consumers some perspective so they can recognize and avoid several surprisingly common failures in critical reasoning that are the enemies of fusion center thinking and good intelligence.

Potholes and Gravel on the Road

One of us worked for a few years as an English-language teacher in China. Every once in a while, the teaching staff would receive feedback that parts of our lessons were incorrect. They would respond with some variation of the statement, “I’ve been speaking English since I was two years old. Why do I need to be told how to do it right?”

This kind of framing raises its invisible head whenever analysts attempt to think critically about events and trends and to draw conclusions. They’ve been thinking about the world since they could first tell colors and shapes apart. Why do they need to be told how to *think*? Yet false confidence about thinking processes often distorts apparently solid analysis.

To clarify our definitions, a **logical fallacy** is a mistake of reasoning that is wrong *inherently* — it breaks a basic rule of logic. A **cognitive bias** is a tendency to think in a certain way that can lead to incorrect conclusions.

If we imagine the route to a sound conclusion as a roadway, then fallacies are like potholes: bad for the car every time. Biases are like gravel on the road: usually not a menace if you are paying attention, but capable of causing the car to veer off the road if you take your hands off the wheel.

Unfortunately, cyber intelligence analysts must often assemble reports under tight deadlines. It is not unusual for business leaders or department heads to demand intelligence

from security teams within a few hours of an internal incident or highly publicized cyberattack. Rushing is extremely dangerous for analysis. The fallacies and biases we discuss below are much more likely to appear in situations where analysts are forced to provide expedited assessments. We strongly recommend giving them time to create research with less urgency and with more rounds of feedback whenever possible. Also, make sure they have the tools to recognize and avoid these human but harmful tendencies.

Logical Fallacies

Since logical fallacies are so much more destructive to good analysis than cognitive biases, publishers of cybersecurity research typically have stringent editorial processes to eliminate them. However, these processes are not perfect, and logical fallacies show up even more often in situations where intelligence is less rigorously reviewed — e.g., in the work of internal security teams.

We think the following questions are useful in identifying gaps and failures in security intelligence caused by logical fallacies.

“Are you sure that your explanation isn’t based on a coincidence?”

“The invalid assumption that correlation implies cause is probably among the two or three most serious and common errors of human reasoning,” said Stephen Jay Gould in his book, *The Mismeasure of Man*.²⁴ This fallacy is the idea that because two things are associated — whether because they occur at the same time, occur in a sequence, have similar attributes, or have been linked before — they are *causally* related.

One of the better examples in a cybersecurity context occurred when a research group linked a hacktivist attack to a nation-state sponsored group merely because both groups happened to have compromised an identical target. In 2016, almost certainly in reaction to territorial conflicts in the South China Sea, the Chinese hacktivist group 1937CN defaced screens at multiple Vietnamese airports with messages supportive of China and

24 <https://www.amazon.com/Mismeasure-Man-Revised-Expanded/dp/0393314251/>

offensive toward Vietnam and the Philippines. After the attack, a Vietnamese security company released a report claiming that 1937CN had been working with the Chinese state-sponsored group APT27 because of indicators that both groups had compromised Vietnamese airport systems at the time of the attack.

This report is no longer public, but even at the time the links between 1937CN and APT27 were spurious and have since invited criticism, such as in a [review of attribution mistakes](#) by FireEye analyst Sarah Jones in 2019.

The original researchers' mistake here was believing correlation equals causation. Their mistake was neither unique nor uncommon. The same fallacy can lead analysts to suggest that:

- ... a Russian APT's intrusion at a French manufacturer must be due to France's support for Ukraine in its war against Russia's invasion
- ... a sudden spike in North Korean crypto-theft campaigns must be due to a rise in the value of Bitcoin
- ... a rise in bot-amplified information operations must be the reason why a particular candidate was or was not elected

The key problem across all those examples lies in the phrase "*must be*." Analysts may have defensible hunches that European support for Ukraine drives Russian APT targeting, or that Bitcoin prices spur cryptocurrency theft, or that bot networks influence voters' information landscape and thereby their choices. Because such hunches quickly come to mind and seem plausible, analysts often accept them as the best possible answer instead of digging further or entertaining other hypotheses.

"Are you assuming that X caused Y this time because X has caused Y in the past?"

The logical fallacy of "affirming the consequent" (also called "the confusion of necessity and sufficiency") is based on the assumption that because X *can* cause Y, it always *must be* a cause or the sole cause of Y. Put another way, this is the mistake of imagining that X *necessarily* causes Y just because X has caused Y in the past. It ignores the fact that other causes may exist.

It is unnerving how often this fallacy shows up in analyses that we've read, from first-round reviews all the way to public reports. Consider the following example, in which both cyber threat and geopolitical analysis are called on to make an assessment:

- Four months after Russia's invasion of Ukraine, a security researcher discovers that a previously unknown piece of malware has compromised the networks of a Ukrainian cloud services provider. The malware has predominantly acted as an infostealer, grabbing as much sensitive information as it can about the provider's clients.
- An analyst looks at the victim and the actions of the malware and concludes that this must be part of a larger Russian effort to infiltrate and destabilize Ukrainian networking infrastructure.

In this case, the analyst has committed the fallacy of affirming the consequent. It is a reasonable premise that if Russia engages in kinetic warfare against Ukraine, it will target Ukrainian companies to gather data and disrupt society. However, it is *not* sound reasoning to deduce that, if a Ukrainian company is targeted with new malware, Russia is responsible. There are hundreds of nation-state intruders and cybercriminals who make use of a cloud provider's data, and while Russia is certainly a *plausible* attacker, without confirmatory data it is by no means the *only likely* attacker.

“Are you just making this argument because someone else said it?”

The fallacy of argument from authority is exactly what it sounds like: accepting a premise or argument purely based on respect for its source. The simple version is to justify an assessment with, “It's true because X said so.”

Argument from authority is pernicious for multiple reasons. First, it is simply logically *invalid* – the conclusions published by domain authorities should only be accepted if they are supported by strong evidence and solid reasoning. In the second place, even sources that are typically reliable can have blind spots because of their perspective, data, reasoning, or prejudices.

We have most often seen argument from authority occur in cyber threat intelligence when analysts fail to notice the point at which a report switches from *observation* to *assessment*. The mistake happens like this:

- Security Vendor Z releases research about a new APT group targeting nonprofit organizations with malware disguised as fake video conferencing software. Vendor Z provides several tables of IOCs associated with the campaign, including fake downloader domains, IP addresses, malware hashes, etc.
- At the end of the report, Vendor Z speculates that the most likely reason for the group's targeting of nonprofits is to gather intelligence on an upcoming diplomatic conference.
- An analyst reviewing this report accepts the validity of the IOCs and is then primed to trust that the rest of the report has been equally well vetted and confirmed. As a result, the analyst takes the report's speculation about the APT group's motivation at face value and repeats it (paraphrased) without seeking any confirmation.

To be clear, in intelligence writing, analysts can and should separate the concept of a *trusted source* from the fallacy of argument from authority. No analyst can independently verify the evidence provided by every report, vendor, agency, and spokesperson. Nor is it practical to attach caveats and hedging language to every piece of data obtained from a source. Rather, analysts can be explicit about the level of trust they have in sources and methods and how it affects the confidence they have in their own conclusions.

“Are you saying something will continue for a while just because it’s still happening?”

The “never-ending story” is a failure of logical reasoning that is apt to show up time after time in analyst predictions about the future of items under investigation. It is a fallacy of prediction that assumes because something is *currently* happening, it will *continue* to happen without change or diminution into an indeterminate future.

The easiest way to isolate the “never-ending story” is to look for the phrase, “We expect [X] to continue.” In many cases, this kind of statement occurs without justification or context; the author simply asks the reader to trust that the trendline of a current pattern is potentially infinite. Some recent examples in cybersecurity are, “We expect ransomware to continue to be a threat to organizations worldwide,” or “We expect that state-sponsored threat actors will continue to target government agencies for espionage of diplomatic information.”

Both of the above statements, by the way, can be accurate but still *illogical* in their production in the sense that their conclusions are merely presented as true without any confirmatory evidence. As a result, the best way for analysts to escape this fallacy is to ask why current activity will persist. Going back to the ransomware example, there are numerous reasons why an analyst could justifiably expect ransomware groups to continue as a threat to organizations, including:

- The ongoing profitability of attacks, which have led to record-setting payments to ransomware groups year over year
- The very low likelihood that ransomware groups’ leaders will face justice, because they operate within countries willing to turn a blind eye to activity as long as it does not impact them

A good question to ask when presented with a “never-ending” trendline is: “What would be most likely to change this trend?” Further, if a situation truly is unlikely to change in the foreseeable future, then it is probably worth considering as the *foundation* for a much more useful assessment. For example, if ransomware is liable to be a fundamental threat to information security for the next several years, then analysts should shift their focus to assessments of the most likely *tactics* and *targets* of ransomware groups in different contexts.

Cognitive Biases

We are heavily indebted to the late Daniel Kahneman and his magisterial book, *Thinking, Fast and Slow*,²⁵ for most of the fundamental thinking about cognitive biases. His ideas are useful to any analytical discipline, but they have unique applications within cyber threat intelligence.

“What would we see most often in the threat landscape if we weren’t focused on the new and interesting?”

One statistical concept can sharpen an analyst’s assessments in every context and for every report: **base rates**. We would put it among the top three concepts that every intelligence analyst should be able to explain and use.

The standard definition of a base rate is the percentage of a population that exhibits a certain trait or characteristic. The trick is making sure you are considering the right population, and not just one segment of it.

For the sake of discussion, suppose that 95% of all the cats in the world are black. In that case, the base rate for black cats is .95. Imagine that you see a group of stray cats emerging from a corner to cross the road. The first three are orange. You may be tempted to predict that the next one will be orange as well. However, your knowledge of base rates would tell you that the odds are strongly in favor of the next cat being black, regardless of any other factors.

That might seem like a silly example, but this bias can have much more substantive consequences in other contexts. It is often one of the primary reasons that we have seen analysts over- or under-emphasize threats. They get caught up in “tunnel vision” looking at a subcategory of malicious activity or geopolitical motivation or technological capacity without considering those in the context of larger, and more statistically relevant, data sets. The way that this shows up most often is as a result of analysts falling victim to the availability heuristic (in which the mind chooses the most recent or interesting data to focus on) but we have seen it appear in other contexts and as a result of other biases as well.

25 <https://www.amazon.com/Thinking-Penguin-Press-Non-Fiction-Paperback/dp/0111573742/>

What does this look like in practice? Here are some examples we have seen across thousands of intelligence reports:

- In response to a question like “What is the threat to our business from APTs?” analysts tend to put blinders on and consider the question from the perspective of a world in which the *only* threat to a business is from APTs. The ability to pull back and consider that businesses face a host of threats from many different cyberattackers is crucial for helping analysts properly estimate the likelihood and severity of an APT targeting an organization.
- In response to a question like “Why is country X targeting our industry?” analysts (especially junior analysts) tend to gravitate towards the explanation that has most recently been offered in cybersecurity reporting. A better approach, one taking account of base rates, is to enumerate *all possible* reasons why a country might target an industry and then run a comparative analysis of the likelihoods of each of those categories of possibility (which should add up to 100%).
- In response to a question like “When will generative AI be able to find zero-day vulnerabilities on its own?” analysts will be tempted to focus entirely on the capabilities of AI as implemented in red team or criminal campaigns. However, the better way to think about AI capabilities with respect to novel datasets is to look at how AI is being used across all possible use cases, cybersecurity-related and otherwise, and to compare those against the necessary conditions for zero-day identification.

“What data are we missing that would most change our perception of the threat landscape?”

“What you see is all there is” is a cognitive bias toward only assessing evidence or data that someone has in front of them. For example, imagine that you were tracking the news in early 2024 and a security vendor identified an intrusion campaign against a U.S. energy company. They attributed the attack to a China-nexus APT group. Based on recently released research

about Volt Typhoon's targeting of energy companies, it would be easy to assume that this new campaign was part of similar repositioning efforts for future disruptive activity.

However, there are other potential pieces of information that might challenge this assumption. What if data was released subsequently showing that the group only targeted one energy company, but targeted more than 30 video gaming companies? What if subsequent data showed clear intent to steal certain information rather than remain on victim systems?

This bias can affect both junior and expert analysts, although for different reasons. Junior analysts are likely to lack enough experience to understand how limited their available data is. On the other hand, senior analysts are likely to succumb to thinking that every problem is a nail to their hammer of expertise: i.e., because they have so much insight into a single field, they tend to ignore data or findings associated with other fields.

One of the better mitigations for this bias is simple skepticism. Analysts and business leaders should look at any assessment about threat activity – especially assessments with heightened emotion or description of severe impact – with some healthy doubt. While an initial assessment might remain the best justified based on available evidence, analysts and business leaders should regularly ask, “What other evidence might challenge this assessment, and how easy would it be to find that evidence?”

“Is this the best description of our threat landscape, or just the easiest?”

The “What you see is all there is” bias falls within the larger category of “cognitive ease” biases. Cognitive ease describes the brain's tendency to select the explanation that is most readily available or easiest to describe, whether or not it is the most accurate. Numerous experiments have demonstrated our desire for quick answers.²⁶ Cognitive ease also means we are more likely to accept a plausible-sounding explanation

26 https://www.researchgate.net/publication/366579357_Availability_Heuristic_An_Overview_and_Applications

even if, by being more specific, it is *less* likely. That can seem counterintuitive, but consider this scenario, which plays off the famous “Linda” problem that Kahneman experimented with in the 1980s.

Imagine that a cybercriminal has just been arrested in Europe for creating a malicious botnet that compromises IoT devices. You are presented with three scenarios:

1. The criminal is a Russian national who has a degree in software engineering.
2. The criminal is a Russian national.
3. The criminal has a degree in software engineering.

Based on the short information about the arrest, which one of these scenarios is most likely?

The most likely is either 2 or 3 (we need more data to decide which one). Scenario 1 is definitely less likely than either 2 or 3, because it involves *more things being true at the same time*.

It can be helpful to think of any attribute of an event or item as part of a Venn diagram. In the circle of “Russian nationals” there is a smaller circle of “Russian nationals with degrees in software engineering”. The second circle will *always* be smaller than the first circle, meaning that the chances of something appearing in the second circle will *always* be less than its chances of appearing somewhere in the first.

Intelligence stakeholders can catch this mistake by looking for increasingly long and descriptive assessments. For example, the assessment “It is likely that our business will be targeted in the next year by Russian-speaking cybercriminals operating an IoT botnet that compromises CCTVs and small-office / home-office routers” has a lower probability of being accurate than “It is likely that our business will be targeted by an IoT botnet in the next year.”

“Does this report look good just because it has numbers in it?”

How many malls are there in the United States? Do you think there are more or less than 10,000?

When the question is presented this way, your answer will probably be much closer to 10,000 than is accurate. This is an example of an anchor – an external data point that your mind cannot help but fixate on. Like a physical anchor, such a data point discourages our minds from moving through the full range of possibilities, leading to a less accurate estimate. (For the record, the number of malls in America is [less than 1,200](#) and will probably be in the hundreds within the next few years.)

Anchors are part of a larger cognitive bias called the *availability bias*, under which the brain prioritizes information that is most readily available, whether or not it's most relevant, accurate, or precise. When we are given a number or statistic as the initial way to understand a topic, it can be almost impossible to wrest our brains away from that starting point.

In the context of fusion center intelligence, the availability of *any* statistic can lure analysts into referencing it as long as it supports their argument. For example, imagine that analysts are working on a report about future threats associated with AI. As part of their research, they find an online article that says AI systems are effective at finding 20% more vulnerabilities than traditional scanners. The analysts take this statistic as gospel and assess that malicious AI-enabled tools will be able to overwhelm defenders.

But did the analysts investigate the statistic? When we probe these kinds of figures, nine times out of 10 we discover that they are either based on very small samples or on marketing materials from vendors attempting to sell a solution with a scary statistic.

In addition, even well-researched statistics can mislead intelligence producers and consumers. Every year, the U.S. Federal Bureau of Investigation's Internet Crime Complaint Center (IC3) releases a report on the volume and types of complaints that the FBI receives regarding cyber threat activity. The [IC3 report for 2023](#) stated that the amount of financial losses reported by victims of such activity grew steadily between 2019 and 2023, with the most recent number being \$12.5 billion.

If asked to comment on financial losses worldwide, an analyst might be tempted to start with that IC3 figure and extrapolate a global total using, say, the International Monetary Fund's [breakdown of worldwide GDP percentage](#) by country. The analyst might propose that, since the United States makes up 15% of the global economy, global losses from cyberattacks are likely to be as much as \$85 billion (\$12.5 billion divided by 0.15).

There are many issues with using the IC3's high-fidelity statistic as an anchor (even leaving aside issues with calculating world GDP). We might ask:

- Is the FBI as effective at collecting information on cyberattack losses as authorities in other regions or countries?
- Couldn't the United States be targeted at much higher rates than other countries?
- How are losses being calculated by the FBI or the victims it interviews?
- Do U.S. victims generally have better or worse security measures than victims in other regions, which might cause them to suffer fewer or more losses?

In short, intelligence stakeholders should see numbers as evidence that may or may not be trustworthy, applicable, and comprehensive. The presence of statistics is no guarantee that one assessment is somehow more thoughtful than another that is more qualitative. With Mark Twain, we advise caution against "lies, damn lies, and statistics."

"Are we misinterpreting a trend as a bigger deal because we're not looking at things over the long term?"

Because of the drumbeat cycle of news associated with cybersecurity reporting and the constant evolutions and updates of threat groups, there are constant temptations to misinterpret minor or short-term shifts in activity as evidence of a huge new trend.

Take the [threat from the LockBit ransomware group in August 2022](#), after a fractious interaction with one intended victim, that it would start engaging in DDoS attacks against its other extortion victims. Or remember when, after the SEC increased penalties for businesses that failed to report “material” ransomware incidents, the ransomware group ALPHV [threatened to alert the SEC](#) about their attacks.

In both cases, analysts could have been forgiven for anticipating nasty trends related to extortionary DDoS and ransomware attacks. However, neither trend materialized. LockBit did not add DDoS attacks to its extortionary toolkit on a recurring basis. ALPHV seems to have threatened to alert the SEC only once, while the novelty of the SEC rule change was in the news. Analysts should have thought more about whether LockBit and ALPHV would need to expand their operations to support the “overhead” of increased threat activity, and whether expansion aligned with the groups’ streamlined style of operations.

One easy clue for spotting this bias in risk assessments is the word “unprecedented.” Although cybersecurity is a relatively young field, practitioners can look back on decades of experience with diverse types of threats. Calling a trend or event unprecedented should raise red flags, suggesting that the intelligence producer is only aware of high-impact activity in the last few months or years.

Epilogue

Fusion Center Rethinking



A Quick Review

In the introduction, we proposed that the modern organization's security program needs fusion center thinking to make cybersecurity more effective by capturing the influence of non-cyber events.

In the next few chapters, we discussed the details, suggesting best practices for team building, data collection, analytical frameworks, risk assessment, and critical reasoning.

To close the loop in accepted writing style, let's recap the main points:

- The risks and threats to an organization arise from a complex landscape of cultural, political, financial, and technological developments. Protecting against these threats requires the ability to integrate knowledge about non-cyber domains into security intelligence. We call this integrative approach fusion center thinking, because it appears most often in organizations that have a fusion center or processes that simulate one.
- Not all domains are created equal. In a widening circle of relevance, analysts should understand cyber data sources inside and outside their organization, then non-cyber issues within their own business, then external non-cyber domains. Each business will invite unique threats based on its own context.
- The most effective cybersecurity data collection starts with good analysts. Internal data is more important than external, and more data is better than less. Data for different domains lives in increasingly unique sources, so general curiosity and awareness are paramount to avoid tunnel vision.
- Building an effective fusion center requires strong executive sponsorship, given the high costs and inherent organizational friction that attend the creation of a group designed to break down silos between subject matter experts. One of the best recent examples of a fusion center in action is a fraud fusion center, where cyber threat analysts document the TTPs of cybercriminals and fraud analysts identify the best measures to detect and combat these TTPs in action and at scale.
- Predictive analysis is only one of several useful assessments to come out of an interdisciplinary security program. Risk assessors should be comfortable talking about a full range of potential threats to the organization's assets, including low likelihood / high-severity scenarios.
- An essential element of fusion center thinking is good reasoning. To achieve it, analysts need to avoid logical fallacies and cognitive biases, and employ strong, situation-relevant analytical frameworks.

And here are brief descriptions of some research findings that did not find their way into this book. Feel free to contact us if you want to hear more about them.

- Ongoing improvement in security requires constantly stretching to secure increasingly complex environments. In the mid-90s, cybersecurity centered on static malware identification. Into the aughts and 2010s, it transitioned to high-volume logging filtering and hygiene. Over the next 10 years, we must focus on better integration of non-cyber contextual information.
- A recent business trend is the flattening of middle management roles. (Consider Bayer, which laid off [thousands of managers](#) in early 2024 to achieve higher efficiency, with [less than stellar initial results](#).) Regardless of whether it is ultimately profitable, this trend reflects a desire to get rid of the kinds of bureaucratic layers that separate production from strategy and encourage gatekeeping. With downsizing on leaders' minds, it's imperative for security teams to be plugged in to the rest of the company's knowledge and priorities.
- How does generative AI affect the future? It will increasingly put expert-level synthesis in the hands of the average consumer. That said, the more generative AI undermines job security for businesses involved in information work, the more likely it is that those businesses will put their proprietary and valuable data behind paywalls or platforms for market advantage. In other words, analysts should be ready for increased functionality in generative AI to lead to less, rather than more, available data for analysis.
- To be useful, the information in this book needs to be applied regularly to security operations. A more contextually informed cybersecurity team is not a checkbox to be ticked for regulatory safety or competitive peace of mind — it is the appropriate next stage once a cybersecurity program has effectively mastered the whole scope of internal and external cyber-specific data.

- Underlying our recommendation for security teams to talk more with other business groups is the assumption that these different teams can talk to each other productively. In other words, if you want a culture of effective communication between security and other departments, you need a corporate culture of effective communication, which needs to be modeled from the top down.

Cybersecurity used to be easier, in part because the threat landscape used to be simpler. Today, threat actors have had decades to perfect certain TTPs and operational setups. Additionally, as software and hardware and cloud and authentication and the entire network of cyber infrastructure have evolved, the requirements for monitoring them and identifying malicious activity have ballooned as well. As in any field, increasingly sophisticated functionality makes it more challenging to understand what is happening when systems come under attack.

The relative simplicity of cybersecurity in the past also reflected the state of the world. In the early aughts, despite terrorist attacks and devastating wars in the Middle East and elsewhere, the sense of a stable world order was still intact. The global economy had not recently experienced a major recession or a worldwide pandemic. There were fewer government-sponsored organizations and allied individuals plotting to perpetrate theft, disruption, and destruction against adversaries.

Fast forward to now, and the world has become far less secure, particularly over the last five years. The security of one's home and paycheck are threatened by inflation and increasingly inaccessible staples. Entire professions face the threat of profound AI disruption, even if much of that threat is hype. The security of political systems in many countries has degraded due to broken civic trust, hyperpartisanship, and information fakery. The security of life and limb is under direct attack now or is likely to be so soon in countries in Europe, the Middle East, and East Asia.

With all these upheavals, is it surprising that organizations of all kinds have attracted more and increasingly dangerous cyberattackers? An unstable world naturally encourages cyber threat actors of multiple stripes: hacktivists have more to protest, criminals see more chances to steal, and nation-states want more information and advantage.

Understanding the many non-cyber domains that influence cyber is not just about anticipating threats better to create a default secure landscape for our collective use of an incredible suite of connective technology. This understanding should also prompt us, in the roles and influence we have, to improve that larger non-cyber world. The better the condition of the average citizen of the world, the less likely we are to see crime and aggression, in cyber or outside of it. For some, that fact may feel uncomfortably close to accountability. That does not make it any less true.

Our hope is that the people responsible for defending against cyber threats become more aware of the drivers for those threats. In this way, they will also become better advocates not just for a more secure world – which everyone deserves – but also for one that is more healthy, more prosperous, more respectful, more stable, and more fair.

About the Authors



David Carver is the director of customer-requested finished intelligence at Recorded Future, the world's largest intelligence company. He has produced or managed threat intelligence for over a decade, tracking hacktivist threats at iSIGHT partners, malign influence operations at FireEye, and larger trends in vulnerability exploitation and malware deployment at Recorded Future.



Levi Gundert is Recorded Future's chief security officer. He has spent 20 years defending networks, arresting international criminals, and uncovering nation-state adversaries. He previously led senior information security functions across technology and financial enterprises. Levi is an author, speaker, blogger, and columnist and a trusted risk advisor to Fortune 500 companies.



Fusion Center Thinking

Extending Cyber Threat Intelligence to Combat Fraud,
Geopolitical Threats, Cybercrime, IP Theft, and More



To be great at cybersecurity today, you need intelligence about all the domains that affect you, including geopolitics, financial fraud, criminal organizations, and unscrupulous competitors. You must obtain input from people knowledgeable about those domains, blend that with existing cyber threat intelligence (CTI), and deliver informed, actionable risk assessments to technical and non-technical managers.

In this book, David Carver and Levi Gundert, threat intelligence experts at Recorded Future, explain fusion center thinking, a proven approach to blending knowledge from multiple domains to help organizations combat fraud, geopolitical threats, cybercrime, theft of intellectual property, and other hazards.

David and Levi explain the basics of fusion center thinking and describe the data sources, staffing, processes, and leadership styles that support it. They outline the critical thinking and risk assessment skills that make it more effective. And they provide examples of how fusion centers and fusion center thinking have proven successful in a variety of real-world situations.

Fusion Center Thinking is a unique guide to concepts and practices that will help every cybersecurity organization improve performance and reduce risk.

