

T H E
ULTIMATE
GUIDE TO
CYBER
THREAT
P R O F I L I N G



TIDAL  CYBER

THREAT-INFORMED DEFENSE

tidalcyber.com

CONTENTS

Introduction: Why This Guide?	3
Why does this guide exist?	4
Who is this guide for?	4
What will you gain from this guide? How is it structured?	5
Why should you take our advice?	6
Chapter 1: The Importance of Threat Quantification	7
Chapter 2: The What & Why of Cyber Threat Profiling	10
Terminology Roulette	11
Cyber Threat Profiling's Value & Strengths	11
Threat Profiling Challenges, Misconceptions, & Limitations	12
Chapter 3: An Achievable (and Repeatable) Approach to Threat Profiling	15
Introducing: Enterprise-Centric Adversary Behavioral Threat Profiling	16
Consider Organizational Context	18
Identify Relevant Threats	24
Quantify Threats	34
Action	41
Chapter 4: Iterating on Your Threat Profile	43
Maturity Opportunities	45
Appendix I	47
Appendix II	49
Cyber Threat Profiling Glossary	51
Endnotes	55
About Tidal Cyber	57



INTRODUCTION:
WHY THIS GUIDE?

WHY DOES THIS GUIDE EXIST?

Recent years have witnessed growing awareness of the benefits offered by a “threat-informed” approach to defense. Most notably, orientation towards the relatively narrower range of possible adversary behaviors provides defenders far more focus than trying to “boil the ocean” of patching each newly reported vulnerability, for example.¹ While growing awareness is an extremely welcome trend, defenders continue to face common practical obstacles to implementing threat-informed defense. Most prominently, too many threats exist in today’s landscape for any single team to reliably track and defend against every one.

The concept of threat profiling offers the potential for threat prioritization, but even when security leaders choose to pursue it, misconceptions over its validity and utility and the lack of a clear and repeatable approach to profiling – as it relates to organization-wide threats – have all hampered its adoption. Even when teams do take steps to prioritize threats, efforts often prolong (in many cases indefinitely) or are impeded by a need for deep intelligence subject matter expertise.

If you are entirely new to the threat profiling discipline and the value of threat prioritization, start with the full introduction presented in [Chapter 1](#). More background on the factors that have traditionally hampered threat profiling’s adoption can be found in [Chapter 2](#). Readers will find the core content of this resource, Tidal’s recommended approach to threat profiling and how this approach addresses existing profiling obstacles, in [Chapter 3](#).

WHO IS THIS GUIDE FOR?

We believe the approach outlined in this guide is practical enough for a wide range of security roles to implement. These include:

- ▶ Security Leadership
- ▶ Cyber Threat Intelligence Analysts
- ▶ Upper and Lower-Tier SOC Analysts/Operators

- ▶ Detection Engineers & Threat Hunters
- ▶ Red/Offensive Security Teamers – Adversary Simulation/Emulation Engineers
- ▶ Purple Teamers
- ▶ Governance, Risk, & Compliance (GRC) Analysts

We have observed cases where practitioners from each type of role had adopted elements (if not large portions) of the approach outlined here (or very similar ones).

WHAT WILL YOU GAIN FROM THIS GUIDE? HOW IS IT STRUCTURED?

Insights from this guide will help you answer the following important questions, which modern security practitioners increasingly face (many are increasingly found in team procedural documentation or “Priority Intelligence Requirements”). Many seem straightforward, but in our experience, analysts, operators, and even leaders often struggle to provide quick answers to them:

- ▶ Which threats matter to our organization?
- ▶ Which threats matter most? (How do I prioritize (rank-order) our list of threats?)
- ▶ How do I take action to address our top-priority threats?

Framed a different way, this guide provides the structure, resources, and tips that allow security practitioners to practically apply multiple threat-related frameworks and methodologies that they might know from academic settings but have struggled to apply effectively in operational settings, including the CIA Triad, the Diamond Model, MITRE ATT&CK®, the OODA Loop, and more.

The guide begins with a Glossary that defines common (and commonly confused) key terms. Next are discussions on the value of threat profiling and challenges and misconceptions that have limited its



TIDAL THREAT PROFILING PRO-TIPS

These sidebars spotlight tips, guidance, or resources that will level-up your threat profiling efforts.

adoption to date. Chapter 3 forms the core of the guide, outlining Tidal’s profiling approach while building a sample profile along the way, complete with immediately applicable resources, tips, and guidance. A large library of relevant resources can be found within the sibling GitHub repository launched alongside this guide: <https://github.com/tidalcyber/cyber-threat-profiling>

WHY SHOULD YOU TAKE OUR ADVICE?

Tidal’s team has decades’ worth of collective experience immersed in the threat-informed defense space. From founding the Center for Threat-Informed Defense², to launching the MITRE Engenuity ATT&CK[®] Evaluations program³ and directly maintaining the ATT&CK knowledge base, to leading threat profiling at a Fortune 20 enterprise (and advising profiling efforts at many other Fortune 100s), our team holds a wealth and variety of insights on practical, effective approaches to threat profiling and threat-informed defense.

Most importantly, our perspectives are informed by countless conversations with defenders supporting organizations of all shapes, sizes, and maturity levels around the world, where we’ve consistently heard practitioners’ challenges with applying threat intelligence. A core Tidal belief is giving back to the community, and we are excited to share this resource in that spirit. A final up-front note: this resource builds upon a growing body of relevant resources graciously shared by many community members – we sincerely thank them for their public contributions and have taken every effort to fully credit and cite others where relevant throughout this guide.





CHAPTER 1:
THE IMPORTANCE OF
THREAT QUANTIFICATION

The need for threat profiling is driven by a single salient fact: too many discrete threats exist in the modern information technology landscape for any single team or even organization to reliably track and sufficiently address (or proactively mitigate) at all times. Recent booms in cyber threat intelligence reporting have likely driven community awareness of various malware and threat actor groups generally. But dramatically lower barriers to entry have almost certainly increased the number of discrete individuals and groups participating in cyber threat activity, while growing attack surfaces provide more opportunities for bad actors to gain initial footholds and move around within compromised networks.

BY THE NUMBERS

Too Many Threats

Mandiant *indicates* it currently tracks 3,500 threat groups in 2023, an increase of 900 from the previous year. The firm also started tracking 588 new malware families in 2022.

In 2023, Microsoft *indicated* that it tracks **300 unique threat actors**, including **160 nation state actors** and **50 ransomware groups**

In 2021, Google's Threat Analysis Group *announced* that it tracks **more than 270 government-sponsored actor groups** associated with **more than 50 countries**

Tidal's analysis of public extortion threats identified **56 ransom groups** that maintained extortion sites in 2022 & 2023

Not Enough Resources

Meanwhile, a December 2022 Neustar *survey* found that 49% of companies did not have sufficient budget to address their cybersecurity needs

As teams are unable to address every threat at all times (to reiterate, this is a fact of modern security operations – every team has resource constraints, some much greater than others) they must take steps to prioritize which threats they dedicate their limited time and resources to addressing. And effective prioritization means that teams ultimately must take steps – even quick initial ones – to quantify which threats matter more or less than one another.⁶No widely adopted, enterprise-centric approach to adversary quantification and prioritization exists in the community today, a critical gap that this resource aims to fill.

Rather than just being an intimidating drag on resources, threat (and related defensive) quantification efforts can help security teams beyond the immediate completion of a profiling project. We expect that this guide will be used most frequently by security teams within private enterprises, who must be able to justify their budgets in financial terms. Quantification supports the generation of metrics and tracking of measurable resulting changes (ideally improvements) in security over time, an aspect that will aid teams (certainly in the private sector but in many ways for public entities too) in justifying their existence and maybe even winning more budget over time.





CHAPTER 2:
THE WHAT & WHY OF
CYBER THREAT PROFILING

TERMINOLOGY ROULETTE

As evidenced by the number of resources referenced throughout this guide, a considerable body of work now exists around the practice of threat profiling. However, after reviewing even a handful of these resources, readers will notice that various terms are often used to describe more or less the same things. The reader isn't mistaken – few consistent, widely adopted definitions exist around the concept of cyber threat profiling as a whole and around many key component and related terms.

To make this book as prescriptive as possible, we provided Tidal's definitions for several key terms in the Glossary. We fully recognize, however, that we are dealing with extremely complex subjects, and being pragmatic, we expect that many teams will have their own variations on these terms – and that's ok. We believe the most important point is that your team picks shared terms and definitions that are most appropriate for your organization and operations, ideally documents them, and remains as consistent as possible when citing moving forward.

*For the purposes of this resource, **Tidal defines Threat Profiling as:** A structured, repeatable process for determining relevant, prioritized cyber threats (adversaries, malware, & associated attack techniques), based on quantifiable evidence.*

CYBER THREAT PROFILING'S VALUE & STRENGTHS

The value of cyber threat profiling is encapsulated within the components of the definition above. Ultimately, the practice of profiling enables organizations to achieve quantification, which enables evidence-based prioritization, which allows addressing the threats that matter most in a timely (or even proactive) manner.

The **structure** provided by a properly developed threat profiling practice means that any member of a

KEY BENEFITS OF CYBER THREAT PROFILING

Structure: Reduces bias

Repeatable: Practical enough to refresh at annual, twice annual, quarterly, or more frequent intervals

Relevant: Don't waste resources on threats that don't matter to the organization

Evidence-based: Enables clear focus, and de-escalation of would-be fires

Proactive: Structure enables identification (and advanced reinforcement/validation of) relevant threat without direct observation

given team should be able to repeat the exercise and achieve relatively consistent results. Ultimately, this benefit *reduces bias* within the results. Again, cyber threat profiling involves complex subjects, so any efforts to limit inherent human analytical bias during the process are useful for generating the most accurate results possible. (We believe that the approach and guidance provided here is practical enough that virtually any security persona, not just intelligence analysts who most often perform these tasks now, can complete the exercise with confidently accurate results.)

Repeatable means that the process is practical enough to be conducted again at future time intervals appropriate to account for the pace of modern adversary TTP evolution, which continues to increase.⁷ Bandwidth- and resources- permitting, this typically means annually, twice annually, quarterly, or in some cases even more often (for many organizations, there will be an upper limit where the value of returns diminishes to the point of being negligible).

Threat profiling is inherently designed to surface threats that are **relevant** to the subject organization. With resources persistently limited for most security teams, no analyst or operator time or effort should be spent addressing a threat that will not likely be encountered (or significantly impact) the organization.

Evidence-based prioritization provides security teams clearer focus, enabling them to justify to leadership that they are attending to threats that really matter. Equally important, it allows de-escalation of would-be “fires” that plague many security teams across the industry. In this sense, we truly believe that profiling can even help address persistent burnout in the industry, contributing to a healthier and more sustainable workforce.⁸

Finally, we believe threat profiling enables organizations to address threats more **proactively**, with confidence (via evidence). The structured nature of the approach outlined here gives teams the ability to identify threats that could or likely would impact their organizations, even if they don't have direct observations of associated activity yet. With proper quantification and prioritization, the team is then able to reinforce and validate their defenses around those threats *ahead of time*.

THREAT PROFILING CHALLENGES, MISCONCEPTIONS, & LIMITATIONS

Several challenges and misconceptions have limited wider adoption of the cyber threat profiling practice to date.

The most notable is a practical one – over the course of many conversations with cyber threat practitioners across the global security community, we have observed clearly that no widely

adopted approach exists to support them in answering common, (seemingly) straightforward questions in a timely manner: “*which threats matter most to our organization, and what can we do about it?*” Several frameworks and methodologies do exist that touch on aspects of the profiling approach outlined here, but common shortcomings have limited their wide adoption.

To be clear, this is not a criticism of these resources – each was created to fulfill certain needs at the time they originated, and many of them indeed informed development of our approach.

FACTOR	LIMITATION
Defensive Scope/Coverage	Asset- or system-centric
Threat Scope/Coverage	Focus on high-level threat categories or scenarios
Complexity	Lengthy, usually require SME input

Table 1: Limitations of existing threat profiling/modeling frameworks & methodologies

Appendix I contains a list of several most-related resources for awareness. We encourage review of this background material, and if you find that certain elements of these frameworks meet your team’s immediate needs, we encourage you to incorporate them into your profiling efforts!

Existing approaches generally fall short in at least one of three ways. First, many existing methodologies support surfacing threats to individual assets or collections of them (systems), but not ones facing an organization as a whole. Others may cover broad threat categories or scenarios, like “denial of service” or “insider threat”, but they fail to detail how those scenarios might actually be conducted, prohibiting translation into relevant defensive capabilities. Finally, most of the frameworks are complex, often requiring granular subject matter expertise to complete necessary information inputs, limiting their repeatability (if they are even able to be completed once).

On top of these practical challenges, misconceptions continue to limit wider adoption of structured threat profiling efforts. The idea of a “threat profile” is not entirely obscure – a search for the term on social media will return many results, most of which are typically sarcastic in tone. Perhaps the powerful potential value of threat profiling makes it seem like an “easy button”, turning it into yet another security buzzword.⁹

Furthermore, the complexity of existing methodologies likely contributes to perceptions that threat profiling can’t be accomplished by the large majority of teams, discouraging adoption. In a similar vein, we have observed the pursuit of near-perfect profiling input data deter or impede profiling efforts, usually related to concerns that not enough data exist to draw meaningful insights around threats unique to organizations of certain types or, in particular,

geographies. While this may have been the case even a few years ago, we believe the quality and quantity of data has reached a point where it should no longer preclude threat profiling efforts.




Figure 1: This chart illustrates the dramatic growth in public, ATT&CK-mapped intelligence reporting, both in absolute terms and relative to the volume of CTI reporting generally (the data derives from a large sampling of threat reporting that Tidal collected & processed). As the volume of TTP- and adversary/victim-mapped intelligence hits critical mass, practitioners are able to derive meaningful insights for threat profiling purposes more regularly.

A final note on profiling limitations: We want to stress that your profile, as defined through the approach outlined next, is often a starting point for further, iterative research and defensive work. Prioritization involves identifying “top” threats, but that doesn’t mean a team should never think about entities lower on their list. Teams *must* start somewhere, and the subset at the top is recommended, but they should ideally continue working down that list as resources and bandwidth allow. The list should also be refreshed, at least occasionally, to confirm that lower-ranked threats actually should remain at that place in the overall order.



CHAPTER 3:
AN ACHIEVABLE (AND
REPEATABLE) APPROACH
TO THREAT PROFILING

Now that we've highlighted profiling's value and addressed common misconceptions, we will spend the remainder of this guide detailing Tidal's approach to threat profiling.¹⁰ Along the way, we will build a sample profile (involving a representative organization but using actual threat intelligence) and spotlight useful resources, tips, and guidance you can immediately implement in your own profiling efforts.

FACTOR	LIMITATION	TIDAL'S THREAT PROFILING APPROACH
Defensive Scope/Coverage	Asset- or system-centric	Enterprise (Organization)-centric
Threat Scope/Coverage	Focus on high-level threat categories or scenarios	Focus on adversaries supports progressive pivoting from organizational context to identification of relevant threats and their capabilities & behaviors, and ultimately to relevant defenses Surface granular adversarial behaviors that align with discrete defensive capabilities
Complexity	Lengthy, usually require SME input	Can be completed by staff with varied skill levels and across team roles/disciplines

Table 2: This table expands on Table 1, highlighting where Tidal's profiling approach addresses challenges or limitations in existing profiling frameworks & methodologies.

The formal label we've applied to Tidal's threat profiling approach is "Enterprise-Centric Adversary Behavioral Threat Profiling". Figure 3 outlines its scope and benefits relative to existing approaches. Our approach directly addresses key challenges and limitations of existing frameworks and methodologies.¹¹

INTRODUCING: ENTERPRISE-CENTRIC ADVERSARY BEHAVIORAL THREAT PROFILING

GOALS

While our profiling approach offers several key benefits, it is worth reiterating that we are ultimately still working with complex subject material: a massive universe of often advanced (and evasive) adversaries that typically have a wide range of software and discrete behaviors at their disposal. The rest of this guide seeks to arm practitioners – in various role types and across skill levels – with practical guidance and supporting resources that enable them to

translate complex topics like Motivation, Intent, Capability, and Quantification into practical leads for profiling research, and ultimately to generate an evidence-based shortlist of relevant, prioritized threats primed for defensive action.

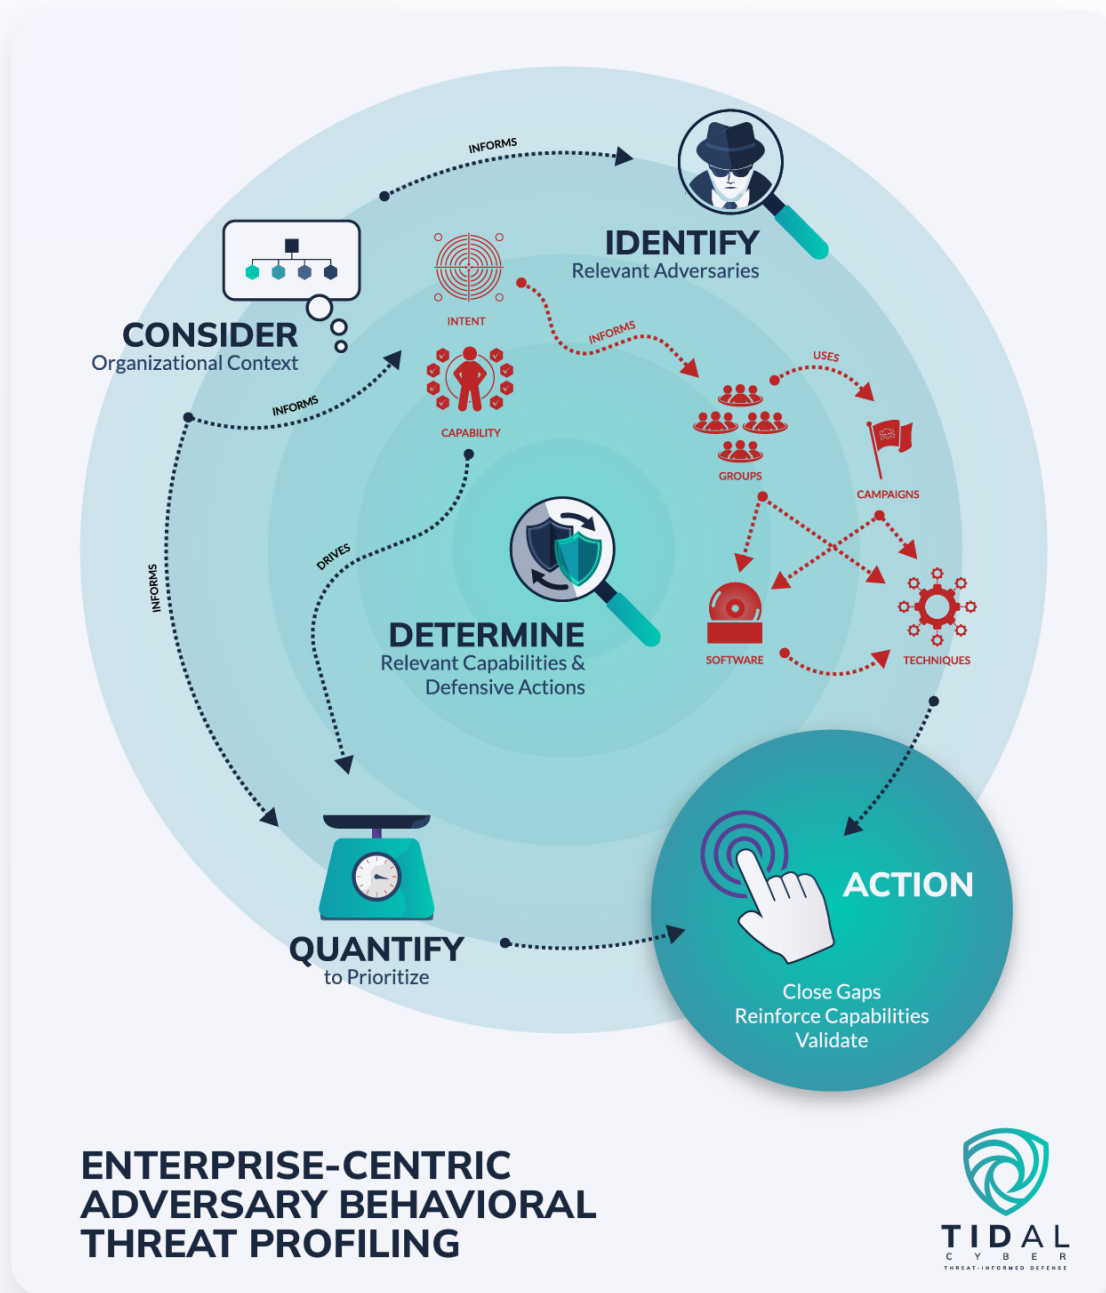


Figure 2: The key workflow and input elements of Tidal's approach to threat profiling.

WORKFLOW

The key distinct elements of our approach include (visualized in Figure 3):

- ▶ **Consider Organizational Context** (p.22): Determine a few logical characteristics most unique to your organization, which informs the general types of adversaries that might impact it.
- ▶ **Identify Relevant Threats** (p. 28): The core of the profiling exercise, comprised of two component phases:
 - ▷ Identify Relevant Adversaries (p. 29): Pivoting on unique organizational factors, surface a list of adversaries that are likely most relevant to your organization based on their motivation or observed proximity.
 - ▷ Identify Relevant Capabilities (& Defensive Alignment) (p. 35): Pivot to capabilities associated with identified adversaries (and/or surface discrete capabilities). Align identified threats to ATT&CK behaviors to unlock further alignment with defensive capabilities.
- ▶ **Quantify Threats** (p. 38): Measure threats according to factors including Proximity/Intent, Capability, density, and organizational priority weightings. Prioritize (rank order) based on relative final weighting scores.
- ▶ **Action** (p. 45): Since defensive action involving behaviors can be resource-intensive, prioritize next steps (often reinforcement, new defensive deployments, or validation) according to threat & organizational priorities, existing defenses, and identified gaps.

THREAT PROFILING PRO-TIP: DOCUMENTATION

Exactly where & how should you document your threat profile? The output of a profiling exercise is ultimately a (prioritized) list of threats (Groups, Software, Campaigns, and especially adversary techniques), so any tool or software that allows you to build & update this list can suffice. In theory, something as simple as notetaking or word processor software could work, although spreadsheet software, especially one that supports simple calculations, will save significant time & effort. As your profiling practices mature (especially as you return to update or maintain your profile on an increasingly regular basis), tools that support further automation, such as scripts, and/or dashboarding software, are highly suggested.

CONSIDER ORGANIZATIONAL CONTEXT

MOTIVATIONAL ALIGNMENT: APPLYING THE DIAMOND MODEL

Before conducting any research queries, we advise first considering the broad types of adversaries that might threaten the organization, which helps orient, validate, filter, and

supplement later research workflows and drive relevance throughout all subsequent phases. Adversaries can be categorized into as few as three buckets based on their Motivations – their

ADVERSARY MOTIVATION	ADVERSARY OBJECTIVES	DISPROPORTIONATELY IMPACTED SECTORS	KEY RELEVANT TACTICS
Espionage	Extract intellectual property from a target	Government, Defense, Technology, Healthcare/ Public Health, NGOs	Persistence, Discovery, Collection, Exfiltration
Financial Gain	Extract Money (or otherwise profit) from a target	E-Commerce, Retail, Hospitality, Travel/Tourism, Payment Processing	Collection, Exfiltration
Destruction	Cause damage or disruption to a target	Production, Transportation, Supply Chains	Impact

Figure 3: Key categories of adversary motivation.

distinct goals during an attack (see Figure 3). While the concept of adversarial motivation might sound like a topic only for intelligence professionals, analysis through the lens of a popular framework, the **Diamond Model**¹², makes it practical for many role types to complete this first phase of the profiling exercise.

A strong threat profile ultimately starts with introspection. Applying the Diamond Model for threat profiling involves orienting to the lower “Victim” node, which represents the subject of the profile. Discovering potential adversary Motivation is then as simple as considering which of the organization’s features represent reasons an adversary might target (or even indiscriminately attack) it.

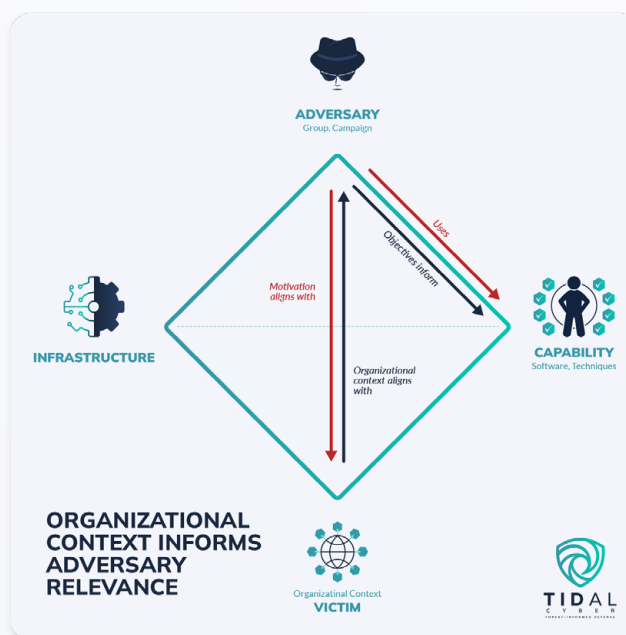


Figure 4: A visual representation of how organizational context informs surfacing relevant adversaries.

Realistically, most organizations in today’s diversified and interconnected business climate will have at least some exposure to adversaries that display each of the three main Motivation types (and many adversaries will exhibit elements of multiple categories). We therefore

recommend focusing on your organization's *most unique features* – the ones that most distinguish it from other entities within or outside its industry sector.

DOCUMENTING ORGANIZATIONAL CONTEXT

We can build more structure around this profiling phase without too much extra effort, subject matter expertise, or even internal organizational knowledge. Since modern enterprises often have large and complex physical and technological footprints, a structured approach provides more confidence that one has captured sufficient relevant context features, but we also don't want to impede the nascent profiling effort with weeks' or more worth of internal discovery efforts.¹³

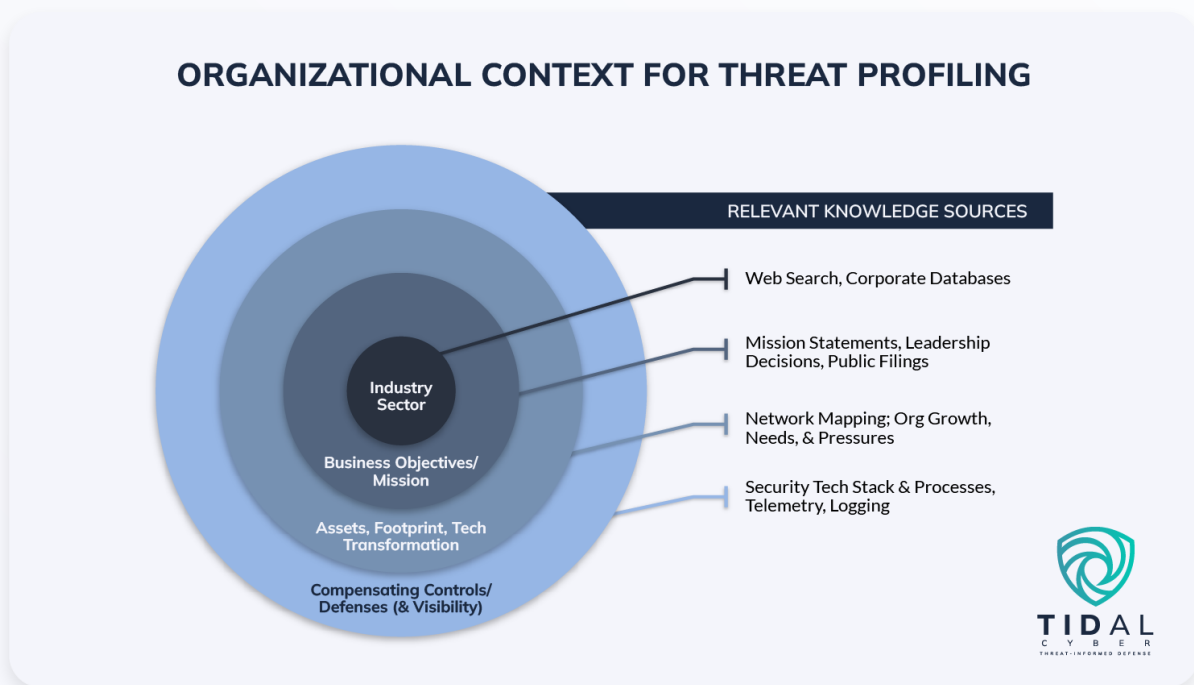


Figure 5: Key information sources for surfacing organizational context details used to inform later phases of a profiling exercise

Unfortunately, our research into existing frameworks and case studies revealed virtually no standardized approach for generating this organizational context relevant for profiling, although some resources do exist that can help inform it without needing to heavily rely on other teams or knowledge bases. A list of them, along with notes on additional nuances to consider, is provided below, sorted by the amount of internal knowledge generally required to provide a meaningful response:

▶ Sector/Industry/Vertical

- ▶ Many organizations can be categorized into multiple sectors. For example, an airport possesses characteristics of the **aviation/airline, general transportation, critical infrastructure**, and even **food & beverage** sectors, while a global hotel chain likely has elements of **travel, tourism, hospitality, e-commerce**, and various other **point-of-sale-related factors**.

▶ Business Objectives/Mission/Functions

- ▶ For public companies in the United States, financial & regulatory filings, mainly those filed with the U.S. Securities & Exchange Commission and especially Form 10-K filings (annual reports), can provide a wealth of information around a wide range of business-related risks, including those relevant to cyber adversarial profiling (see examples below).¹⁴ Internal-facing business updates or reports, especially those that touch on business and especially digital/technology trends in the enterprise, will often be directly relevant.

▶ Assets, Footprint, & Technology Transformation

- ▶ This loosely translates into well-known organizational “People, Process, & Technology” terms.
- ▶ For enterprise threat profiling purposes, we mainly mean “assets” from an impact perspective and less from a vulnerability management standpoint. For example, an organization with a large volume of cloud or container assets may be especially exposed to cloud-focused actors like TeamTNT and resource hijacking attacks like cryptomining¹⁵



**THREAT PROFILING PRO-TIP:
IMPORTANCE OF TECH
TRANSFORMATION FACTORS**

The list of organizational context features should be periodically refreshed in order to check potential bias or assumptions. This is especially true for newcomers to a given organization, but also veterans. For example, it might surprise you to know how much a major home improvement retailer – well-known for its physical shopping centers – emphasized digital platforms in its 2023 annual report, including a reference to the rollout of next-generation digital phones to each of its 471,000+ global associates. A development like this is sure to influence the organizational attack surface and, by extension, potential exposures, vulnerabilities, and ultimately adversaries keen to exploit them.

- ▷ Here we mainly mean the type and geographic footprint of key technology and physical/people assets, which may point to adversarial motivation, and less a comprehensive asset inventory.
- ▷ Consider how planned or unexpected changes in physical footprint, working arrangements (e.g. remote work), technology, and even budget/finances might influence factors relevant for your threat profile.
- ▶ Security/Defenses
 - ▷ Now is a good time to take an initial inventory of compensating controls and defenses, especially if knowledge or visibility gaps are identified that may take time or reliance on other teams to fill.

CASE STUDY: A REPRESENTATIVE PHARMACEUTICAL MANUFACTURER

Here we launch a case study that we will follow through the rest of this guide in order to demonstrate practical application of the workflows, guidance, resources, & tips we provide throughout.

The sample subject is designed to represent a generic large pharmaceutical manufacturer. The company produces a range of specialized drugs, including vaccines for the COVID-19 virus. The enterprise is headquartered in the United States, has major administrative & production sites there and in Western Europe, and supplies ingredients from around the world, especially East and South Asia. It employs 70,000 people.



We will assume the persona of a security team member who does have a deep CTI background and isn't steeped in research or quantification around adversary threats, to demonstrate our assertion that the practical nature of our approach enables those in a range of roles and with varied experience levels to complete profiling exercises, lowering barriers to entry and driving further adoption of the practice.

How could we begin to generate some organizational context for a pharmaceutical producer? Since it sits within the manufacturing sector generally, a natural starting point likely involves considering potential cyber-related interruptions to *physical production functions*. This directly aligns with the disruption/destruction-focused adversarial motivation category. If desired, we

can expand beyond this single example of context-to-motivation alignment, which will help build a deeper and richer (but still relevant) list of adversaries in the next phase. And we can do this without too much added time or specialized knowledge, by using publicly available information (we used a real Form 10-K annual report for this example).

Figure 6 shows the ultimate results of a very quick and rough organizational context assessment. We used the 10-K report’s “Business Overview” section (the first section in the 160-page report) to quickly identify top business functions for our sample organization (Column A). Additional details drawn from elsewhere in the report (Column B) elicited further detail around relevant factors that can be logically linked to potential adversarial motivations.

A case like a production or supply chain interruption is arguably straightforward to link to the Destruction motivation category. In cases where the relationship might not be clear, another popular model, the **CIA Triad**, can provide helpful structure around your efforts – it may be more straightforward to first link a business function to one of the Triad’s data



Figure 6: The CIA Triad

Confidentiality, Integrity, or Availability nodes, then pivot further to a discrete adversary motivation. For example, the impact of a cyber-related disruption to Research & Development may not be immediately clear, but further context from the 10-K report reveals that the organization specializes in producing “highly differentiated” medicines, including vaccines, whose formulas would threaten Confidentiality if exposed. (Intellectual property exposure furthermore directly relates to the Espionage motivation, while we also deemed it a notable Financial Gain motivator since criminals might especially seek to monetize IP exfiltrated from this organization, given its stated high value.)

Australian cybersecurity authorities recommended a CIA Triad-based approach to jumpstarting profiling efforts in a 2020 report, which suggested rank-ordering business services by a numerical rating according to potential impact to the Triad’s components (unfortunately the report did not offer much specific guidance on how to generate those ratings).¹⁶ While not using “CIA” framing specifically, a 2022 webcast from the Red Canary corporate security team outlined a similar process for kickstarting profiling efforts, beginning with asset (data, systems, and financials) discovery, translating those assets into elements of an “attack surface”, and using the attack surface outline to drive identification of relevant adversaries.¹⁷

For some final flare, Figure 7 includes an “Estimated Financial Impact” column, which we populated by simply bucketing according to breakdowns derived from financial figures provided later in the original Form 10-K report. Since we are just trying surface (and quickly prioritize) some contextual considerations to inform our next steps, a qualitative reordering of the lines in this table (in our case by simply “eyeballing” the cells) is likely sufficient for us to now move on. However, this process could probably be quantified without too much effort, for example by tallying either (or both) of the CIA and Motivation groupings and combining those tallies with a numerical score for the Impact column, perhaps via a weighted average.

Organizational Considerations		Priority Ranking	Effect on CIA if compromised			Adversary Motivations			Est. \$ Impact
Business Function (Annual Report Section)	Analyst Context		C	I	A	Espionage	Financial Gain	Destruction	
Raw Materials	Procure from numerous suppliers worldwide; demand could result in supply constraints	1			✓	✓	✓	✓	High
Research and Development	Deliver "highly differentiated" medicines & vaccines	2	✓			✓	✓		High
Information Technology & Security	"Extensively rely" on sophisticated IT systems (including cloud services). We process, store & transmit large amounts of confidential information, PII & IP	3	✓		✓			✓	High
Collaboration and Co-Promotion Agreements	We collaborate, license, acquire, & invest in external parties (universities, biotech) for R&D	4	✓	✓		✓			Medium
Human Capital	A new, flexible working model enables work to be regularly conducted from home	5	✓		✓	✓			Medium
International Operations	Global operations - we supply to 180+ countries	6			✓			✓	High
Sales and Marketing	We adapted our digital promotional platform to support the shift to hybrid virtual and in-person engagements	7		✓			✓		Medium
Environmental Matters	We are currently remediating environmental contamination from past industrial activity	8			✓			✓	Low

Figure 7: Generating organizational context for a representative pharmaceutical manufacturer.

IDENTIFY RELEVANT THREATS

This research-focused phase represents the core of a threat profiling exercise. As we’ve noted multiple times, **there are more threats than any team can possibly address at all times**, so our goal here is to identify the subset out of the universe of threats that is most relevant to us, in order to drive the highest return on investment possible for allocated defensive resources.

The workflow for this phase is generally guided by the definition of a “threat”:

Threat = Intent x Capability x Opportunity¹⁸

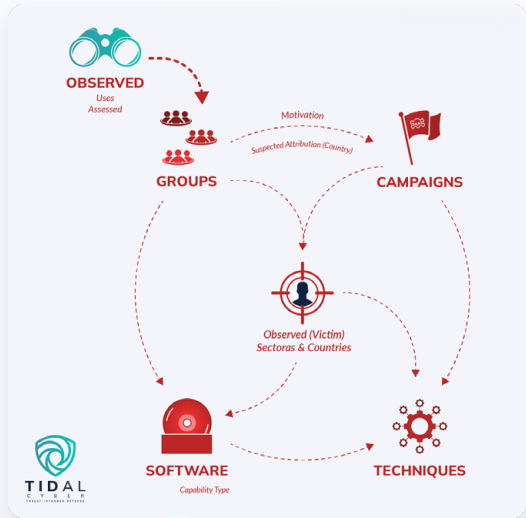


Figure 8: Key types of threat elements and types of metadata commonly associated with each in public or commercial CTI knowledge bases.

We will first surface adversary Groups and Campaigns relevant to your organization and its unique context (sector and/or location), which helps narrow our focus from the immense overall threat landscape to those threats that are most relevant to our specific organization. Sourcing considerations will help us approximate these adversaries' Intent, which will influence quantification efforts later. We will then pivot to (and, where needed, perform additional discovery around) relevant adversarial capabilities, including adversary Software and, critically, discrete behaviors, which we will align directly with our security capabilities in order to optimize future defensive action.

IDENTIFY RELEVANT ADVERSARIES

Intent & Proximity

The Threat equation suggests that, rather than blindly searching for any and all adversaries, we want to especially look for ones with Intent to attack us (without it, we wouldn't consider them a "threat"). But Intent implies assessment of human psychology – ultimately a human adversary is on the other side of an attack, and we're often dealing with actors that are highly skilled, well-resourced, and intentionally evasive. How can we gauge their intent, especially if we don't have a substantial background in intelligence analysis, or state-level intelligence capabilities?

While we rarely have clear evidence pointing to adversaries' ultimate intentions, we can look to the growing body of cyber incident evidence to gauge approximations of them. We believe a critical mass of data now exists, including in the public realm, to enable defenders to begin drawing meaningful profiling insights from it. A practical approach to approximating adversary intent involves looking at evidence of adversaries' proximity to their targets, which we will measure with common attack metadata, including victim industry sectors and locations. We will also lean on proximity tiers to provide approximate Intent scores later.

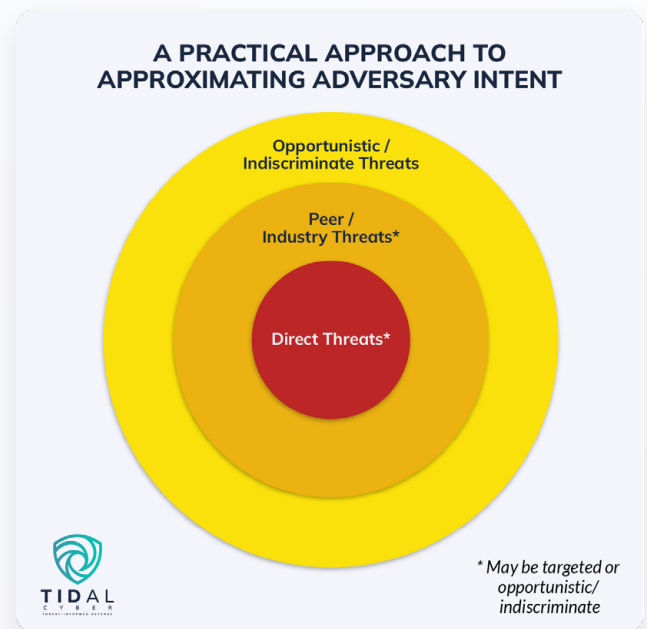


Figure 9: An identified threat's Proximity to the subject organization can be used as a rough estimate of the Intent of the adversary behind it

SURFACING RELEVANT ADVERSARIES

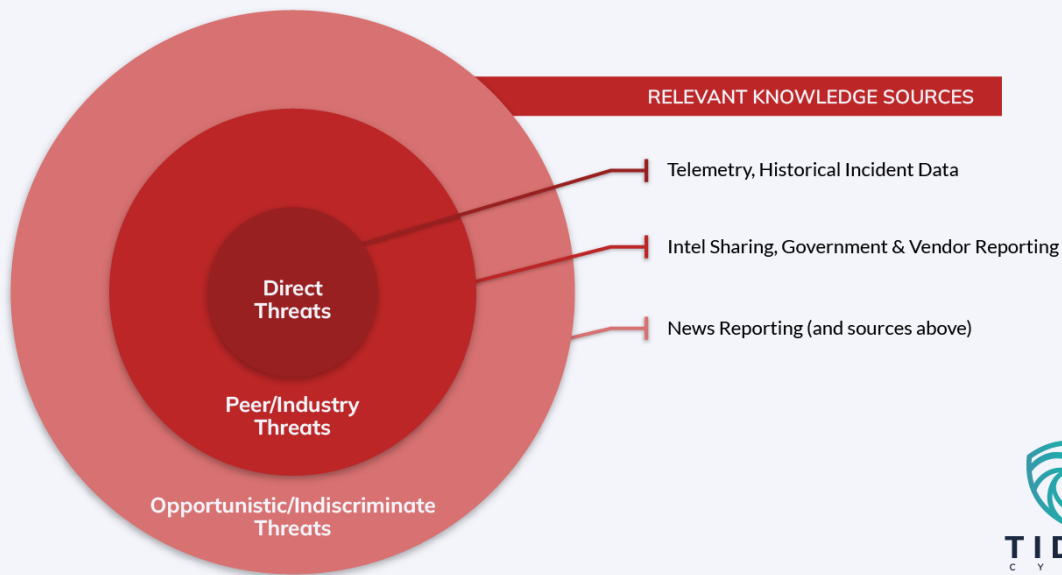


Figure 10: Key information sources for surfacing adversarial intelligence, organized by Proximity bands

Direct Threats

Adversaries known to have impacted your organization should typically be the first ones included in your nascent profile, since direct observation generally provides the most reliable indicator of adversarial intent.¹⁹

Hunters, investigators, or responders may be able to correlate data points from internal telemetry (network and endpoint logs) and attribute observed activity to a particular adversary Group or Campaign. For practical reasons, many organizations will not have the capacity to perform this attribution often. Indeed, we encourage many teams to not necessarily be consumed with attribution-level investigations – due



THREAT PROFILING PRO-TIP: HOW MANY ADVERSARIES IS “RIGHT”?

An important aspect of this phase is surfacing a “manageable” volume of adversaries. The boom in threat intelligence in recent years means that many organizations could finish this phase with a large number of adversaries, rather than only few, which was often the case even a few years ago.

The threshold will vary for virtually every organization, but in our experience, following the approach outlined here, a list of 10-20 Groups and Campaigns (and similar or slightly larger list of Software) is generally more than sufficient. Smaller lists also function adequately.

to ever-present resource constraints, once enterprise security teams have contained and remediated an incident, they must often move quickly to the next one. Fortunately, as we'll show, other reliable sources exist that can be used for profiling beyond just those derived from resource-intensive, attribution-focused investigations. Take caution against adding too many threats entities, which can generate a very large (and potentially unmanageable) number of discrete behaviors, although focus on technique "density" or overlap should still clearly spotlight a subset of top-priority ones.

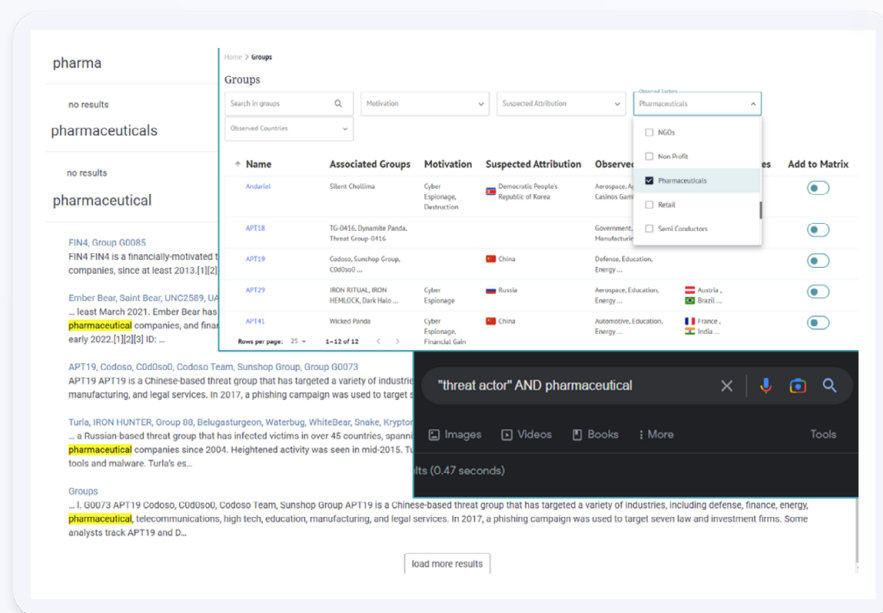
A couple final notes. We have observed that vendor-correlated alerts increasingly include an assessment of adversary attribution, providing another potential source that involves internal telemetry but doesn't require significant natively developed capability. Finally, don't forget to check to see if your organization already has a running registry of historical, attributed incident data, derived from either internal investigations or from vendor assessments.

Proximate Threats

Directly observed threats typically comprise only a portion of an organization's profile, and most organizations will want to look outward to surface additional potential threats. A natural next step is considering threats known to impact other organizations that most resemble your own, including peers or others within your sector, and/or entities with operations in locations matching your own. We typically see that teams surface the largest number of adversary inputs to their threat profile from this bucket.

The boom in intelligence sharing in recent years, across both public and closed sources, has birthed a large body of evidence linking particular adversaries to observed threat activity. Several great sources of adversary intelligence now exist that regularly include metadata around victim sector, location, and/or

Figure 11: Traditional workflows for surfacing proximate threats (pictured at Left and Lower-Right) involve searching unstructured data, which is time-consuming and prone to human and machine error. The structured Group metadata in Tidal's free Community Edition (pictured at Upper-Right) supports quick and accurate searching, saving considerable time & effort.



organization size, including government and national CE/IRT (“computer emergency/incident response team”) advisories, public or commercial vendor reporting, and independent analyses (e.g. incident response or malware analysis/threat research blogs).


We have commonly observed workflows where teams will search across these sources, often via general web or news search queries, for keywords related to their sector or geography (see Figure 11). We find that resources that aggregate adversary metadata *in a structured way* save immense time and effort for profiling purposes. Recognizing this value, the Groups page in Tidal’s free Community Edition (app.tidalcyber.com/groups) provides structured victim sector and location metadata, derived from ATT&CK and many other public sources, for a large number of adversaries, with opportunities to visualize or pivot to other defensive-oriented enrichment around the adversaries and their behaviors. Several other great sources for structured threat metadata, which we see used throughout the community, include:

- ▶ [ETDA/ThaiCERT: Threat Encyclopedia](#)
- ▶ [AlienVault OTX](#)
- ▶ [MISP Threat Actor Galaxy](#)
- ▶ [SecureWorks Cyber Threat Group Profiles](#)
- ▶ [Palo Alto Unit42 Playbooks](#)
- ▶ [CrowdStrike Adversary Industries](#)
- ▶ [APT Groups & Operations \(public Google Sheet\)](#)

Privileged intelligence sharing circles, like those provided via Information Sharing and Coordination (“ISAC”) bodies, represent a key private/close/privileged source of such relevant information. For example, many of these groups will share internal facing metrics around findings reported by their members (e.g. phishing attempts) which may even include more likely attribution since sharing here is governed by TLP restrictions (remember to still use your own independent judgment on shared intelligence’s reliability, though).

We want to emphasize the importance of critical thinking throughout a threat profiling exercise, including during this phase, and discourage overly rigid adherence to the research workflow for surfacing Proximate Threats (or other factors, especially in the quantification phase later, for that matter). While we’ve taken efforts to recommend generally higher-confidence sources for surfacing threats relevant to certain factors like sector and location, there can be times where research surfaces particular threats that simply don’t “make sense”. This is especially true if one is using tools based on machine recognition of industry or geographic identifiers, but it can result from human-tagged results as well. Aggregation resources are extremely powerful for often quickly providing a sizable list of initial results,

THREAT PROFILE DRIVEN BY WHAT



THREAT PROFILE DRIVEN BY WHY

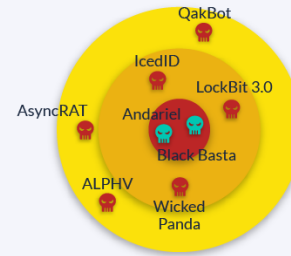


Figure 12: Captures the phenomenon of over-emphasizing any or all threats surfaced via earlier workflows. When we layer on organizational context factors derived via critical thinking, we can narrow down the often-large list of seemingly “relevant” adversaries we have generated so far via our profiling workflows.

but care should be taken to validate that list through the lens of a critical eye, and at least quick supplemental research is often required for threats that you may be less familiar with, to confirm they likely actually possess motivation relative to your organization’s unique contextual characteristics.

Indiscriminate Threats

The previous workflows are designed for filtering the immense volume of threats present in the overall landscape. These flows address major existing obstacles to threat prioritization, but we do want to caution against considering only threats surfaced from this research – in today’s landscape, it is also essential to consider underlying factors (and associated threats) that may be shared among a wide range of entities, including ones that haven’t surfaced in the research flows so far.

The rise of the “as-a-service” megatrend in recent years demonstrates why most organizations should also consider opportunistic and otherwise indiscriminate threats within their profile.²⁰²¹

Many threats today, including highly capable and impactful threats like many ransomware operations, appear to attack almost anyone. They work with or alongside “access brokers” who specialize in gaining initial footholds into a wide variety of networks, often reselling that access to the highest bidder (or working with preferred partners), increasing their range and variety of potential victims. Others perform widespread, often automated scanning campaigns to identify virtually any exposed assets that might be vulnerable to a given exploit (technology-based, or otherwise).

Our report on prioritizing among ransomware-as-a-service operations provides guidance relevant for prioritizing among a variety of indiscriminate threat types.²²

In order to achieve some narrowed focus in the wide world of these potential threats, we recommend leaning on metrics wherever possible, even if the scales do not perfectly align across different reports on discrete threats. For example, data extortion threats made by ransomware groups can be measured and associated with alleged victim size, geography, and sector, allowing us to rank order and identify specific groups that might be most relevant to our organization’s profile. Technical sources like malware sandboxes can also provide quantified indications of when certain threats might be “trending”, another good indicator for raising an indiscriminate threat’s priority level. When all else fails, considering what threats are “in the news” really isn’t a terrible starting place from which to gauge a threat’s potential trendline, although we always encourage further, more rigorous analysis where possible.

Case Study: Identifying Relevant Threats for a Pharmaceutical Manufacturer

Figure 13 shows the results of completing the research workflows for identifying relevant threats. There was a relatively smaller number of directly observed

Adversary or Campaign Name	Proximity Tier	Evidence
Andariel	Direct Threat	Our team attributed with moderate confidence a 2021 incident to this group
Wizard Spider	Direct Threat	Our endpoint vendor has quarantined multiple samples of malware distributed by this group
TA1337	Direct Threat	Our email security vendor blocked phishing emails attributed to this group
Kimsuky	Proximate Threat	Per news reporting, carried out a campaign targeting our largest competitor (our closest peer)
APT42	Proximate Threat	Per our ISAC, targeted another peer org (one that produces significantly different medications but with a similar geographic & supplier footprint)
APT29	Proximate Threat	Per public reporting, linked to several attacks involving large vaccine manufacturers (specific companies not known)
APT9	Proximate Threat	Per public reporting, linked to attacks involving entities in our industry generally
APT41	Proximate Threat	Per public reporting, linked to attacks involving entities in our industry generally
menuPass	Proximate Threat	Per public reporting, linked to a very high and consistent number of attacks in two of our top supplier countries
FIN4	Proximate Threat	Per public reporting, linked to attacks involving entities in our industry generally
WildNeutron	Proximate Threat	Per public reporting, linked to attacks involving entities in our industry generally
QakBot H2 2022-H1 2023 Campaigns	Indiscriminate Threat	Public reports suggest the volume of QakBot infections recently spiked. Attacks were observed in a range of industries, but ISAC members have yet to report cases
Recent RedLine Stealer infection campaigns	Indiscriminate Threat	Public vendor reports indicate a significant rise in RedLine Stealer infections. Recently used lures appear to target the tech & education sectors

Figure 13: The output of the Identify Relevant Adversaries phase for the pharmaceutical manufacturer sample organization. More Proximate Threats were identified (around 30 total) but some were excluded from the image for space considerations.

threats, derived from a mix of internal investigations and vendor attribution. There are relatively many proximate threats (more than included in the image), derived most often from public reporting. We surfaced the majority of these results quickly using the Tidal Community Edition Groups page, supplemented with manual, keyword-based research (e.g. searching pages for the text string “pharma”) involving the following resources mentioned above: ETDA/ThaiCERT: Threat Encyclopedia, MISP Threat Actor Galaxy, SecureWorks Cyber Threat Group Profiles, Palo Alto Unit42 Playbooks. For completeness, we generally recommend including at least a few additional top or trending adversaries or campaigns in the Indiscriminate Threat category, but these examples show types of threats generally worth considering for this phase of a profiling exercise.

DETERMINE RELEVANT CAPABILITIES (AND DEFENSIVE ALIGNMENT)

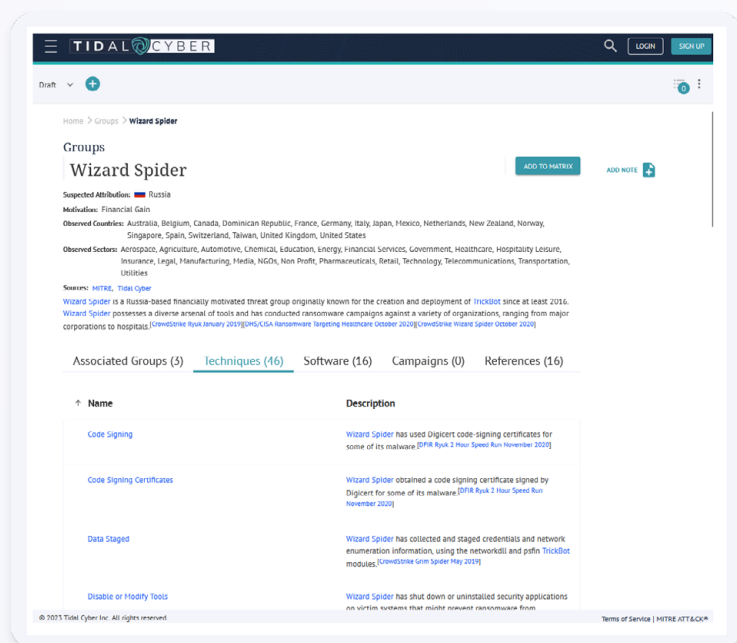


Figure 14: Profile pages in Tidal’s free Community Edition, like this one for a [major cybercriminal group](#), enable instant pivoting from the adversarial level into relevant Capabilities, such as malicious Software and Techniques used by a particular threat.

We recommend leaning on intelligence mapped to MITRE ATT&CK Techniques and Sub-Techniques for the workflows covered in this section, which dramatically streamlines much of the upcoming research effort. More details on the value of using ATT&CK for threat profiling are provided in a dedicated breakout.

Adversary Capabilities

Identifying relevant adversaries helped narrow our focus to just the subset of the entire landscape that is most likely to threaten our organization. But this represents just one piece of the overall threat identification workflow – we must

now discover *how* those adversaries actually carry out their attacks. This phase focuses on pivoting on the knowledge we just surfaced (a shortlist of adversaries and campaigns) to the specific Software and behaviors those adversaries are known to use (and finally, sometimes performing supplemental research to fill in potential research gaps up to this point).

Refer to Figure 15 for a reminder on distinctions and relationships between Adversaries, Campaigns, Software, and Techniques. Certain intelligence sources, including popular public ones, associate capabilities with particular adversaries, with varying degrees of data

structure. Where relevant, the MITRE ATT&CK knowledge base provides both Software (categorized into “Tools” or “Malware”), as well Techniques and Sub-Techniques that map to particular adversaries ((Sub-)Techniques associated with Software’s technical functionality are also provided). Several of the other sources useful for identifying Proximate threats also link adversaries to Software.

Like in the previous workflow, we recommend building a list of the Software and (Sub-) Techniques associated with the adversaries you previously identified (and by extension, any (Sub-Techniques) associated with the newly surfaced Software in your new list). As you build these new lists, we recommend noting where multiple adversaries are associated with the same Software and (Sub-Techniques) – these will form a foundation of the upcoming quantification workflows!

In our experience, adversary-to-Software or -Technique pivoting often forms the main basis for this research phase. However, we highly encourage including capabilities even if a link to a specific adversary is not known. Reasons for this may include current lack of attribution to a defined Group or Campaign, or lack of detailed information around a currently linked group. An instance of the latter case is vendors that track yet-unnamed Groups or Campaigns (often labeled “UNCs” or Uncategorized adversaries) – for example, few concrete attribution details (e.g. an associated origin country) may be publicly disclosed about a given UNC group, but if it is known to use certain Software or behaviors during its operations, and those activities are relevant to your organization, we highly recommend adding those capabilities to your list asynchronously.

You can proactively expand on this workflow by replicating previous workflows and searching resources for indications of Software impacting particular sectors, geographic locations, and/or sizes of business. In our experience, those metadata are not typically tracked as regularly for Software versus Groups and Campaigns, however. To avoid reverting back to the entire canvas of threats and thus adding another large workload to your profiling efforts, we recommend narrowing your research to any categories of Software your organization might already deem as priorities, informed by the organizational context factors surfaced earlier in your profiling efforts. These often include categories like ransomware, offensive security tools, remote access trojans (“RATs”), infostealers, or denial of service, wiper, or initial access threats.

Aligning Adversary & Defensive Capabilities

The definition of threat-informed defense implies that understanding adversary behavior is critical to effective cyber defense – without this, defenders are left to address any and all potential indications of an attack, an approach that is unsustainable amid the immense scope of today’s threat landscape, as spotlighted throughout this guide.

By providing consistent definitions for attacker behaviors (Tactics, Techniques, and Procedures) that are referenced by practitioners throughout the CTI, offensive security, and defender communities, ATT&CK provides a common language used by teams within and across different organizations, at a level of abstraction appropriate for capturing the broad range of potential attacker behaviors and at an appropriate scope and depth to be manageable by frontline defenders.

We view ATT&CK as a foundational component of a strong threat profiling exercise. By focusing on adversarial intelligence aligned with ATT&CK, we can translate – in a straightforward yet accurate way – from the external adversary space into discrete, relevant defensive capabilities. This allows rapid assessment of where current security measures might fall short against threats we care about and where there might be sufficient or even redundant defenses against other attacker techniques. ATT&CK’s wide adoption has contributed to the creation of mappings for other important resources, such log sources and proactive controls and mitigations, broadening the applicability of the knowledge base.

While adversary behaviors by nature change less frequently than the infrastructure used to launch their attacks, we continue to observe a higher pace of adversary TTP adaptation and evolution, often in direct response to improve security posture (a good thing!). This necessitates more consistent intelligence updates around adversary behavior. We are thankful and encouraged to see growing ATT&CK adoption across the defender and threat intelligence communities, which makes this tracking far more efficient!

Case Study: Determine Relevant Capabilities for a Pharmaceutical Manufacturer

Figure 15 shows the truncated result of surfacing relevant capabilities for the sample pharmaceutical manufacturer. We used Tidal Community Edition to pivot from each Group or Campaign in the list of adversaries generated during the last phase (see Figure 13) to surface the Software and the (Sub-)Techniques associated with it, according to the ATT&CK knowledge base (each of the (Sub-)Technique lists and some of the Software lists are also truncated for space considerations).

As the figure shows, supplemental research to surface relevant techniques is often needed. For example, since our sample organization is concerned about QakBot, which notoriously changes its TTPs often, we felt it important to surface techniques observed more recently than those currently provided in the ATT&CK knowledge base, which date through September 2021.

Key CTI sources that most often contain ATT&CK mappings, or provide enough detail to be manually mapped, include:

- ▶ Government advisories
- ▶ Public or commercial vendor threat research & intelligence reports
- ▶ Independent blogs on incident & campaign responses/investigations and malware technical analysis

Also note that we added LockBit 3.0 to the list of capabilities, even though it does not have a single, clearly defined associated adversary group, because it recently carried out an attack on a peer organization. LockBit 3.0, the leading global ransomware in terms of public victim count in 2022, is also not included in the ATT&CK knowledge base, so we added a set of techniques associated with it in public CTI reporting.

While this doesn't appear in Figure 15, remember that you will also want to generate lists of (Sub-)Techniques associated with the Software surfaced in this phase. These can similarly come from pivoting in ATT&CK knowledge base data (a great starting point) and from your own research.

Adversary or Campaign Name	Techniques Used	Software Used
Andariel	T1546.008, T1560.001, T1210, T1110, T1059, T1552.001, T1005, T1005, T1039, T1087.002...	China Chopper, Ngrok, Pay2Key, PsExec
Wizard Spider	T1587.002, T1562.001, T1087.002, T1078.002, T1055.001, T1041, T1048.003, T1210...	TrickBot, AdFind, Bazar, BloodHound, Cobalt Strike, Emotet, Mimikatz, Nltest, PsExec...
menuPass	T1560.001, T1119, T1070.003, T1005, T1039, T1140, T1574.001, T1574.002, T1087.002...	AdFind, certutil, ChChes, Cobalt Strike, esentutil, EviGrab, Impacket, Mimikatz, PsExec...
QakBot H2 2022-H1 2023 Campaigns	T1027.001, T1059.001, T1574.002, T1547.001, T1010, T1140, T1021.005, T1005, T1518...	QakBot
Unknown / Various	T1486, T1489, T1547.001, T1070, T1112, T1548.002, T1078.002, T1574.001...	LockBit 3.0

Used by multiple adversaries
User-defined Techniques

Figure 15: Truncated results after pivoting from the list of discrete adversaries & threats surfaced in the last phase of the exercise, into associated Capabilities (Software & Techniques).

Readers will observe that after pivoting on even a handful of adversaries, a fair amount of overlap in associated Software and (Sub-)Techniques starts to be observed. The scale of this overlap will be used shortly in the upcoming threat quantification workflow.

QUANTIFY THREATS

Fortunately, this phase will lean heavily on previously conducted research, including the proximity tiers used while identifying relevant adversaries, and the “density” or overlap of relevant adversary capabilities (Software and Techniques). For each workflow, we encourage critical thinking when assigning quantification weights, which ultimately serve as approximations of potential threats, rather than following rigid guidelines (including our own recommendations!). Since we are ultimately taking many complex factors into account, a single quantification metric that might be slightly “off” is unlikely to dramatically impact the final results of your threat profile. The quantification guidance provided in each section builds upon multiple existing, fantastic community resources.

GUIDANCE FOR QUANTIFYING COMPLEX THREAT CONCEPTS

The Threat Box model for quantifying threat actor assessments, developed by Andy Piazza, and the National Institute of Standards and Technology (“NIST”) Guide for Conducting Risk Assessments represent foundational studies in adversary threat quantification, and readers will notice overlap with the weighting guidance outlined here.²³²⁴

Our guidance expands upon this foundation, promoting flexibility in analyst judgment based on organizational context, while also accounting for greater depth and specificity in the Capability realm, i.e., the Software and Techniques used by priority adversaries.

For teams performing their first or first few threat profiling exercises, we generally recommend five-point quantification scales. From our many conversations with practitioners, we assess that this scale provides the greatest range of flexibility and variation while still retaining some semblance of practical analyst judgement and implementation. We have observed teams effectively apply scales as narrow as two or three points and as high as 100 points. A higher scale generally requires greater resource commitment for it to be implemented effectively.

Intent

- ▶ Most generally, we recommend weighting adversary intent according to the proximity tiers by which you surfaced each adversary earlier. Direct Threats will generally receive the highest weighting scores, followed by Proximate Threats, followed by Indiscriminate Threats. If you assess that overwhelming evidence exists to move a given adversary’s weighting into a higher or lower level, we encourage you to do so!

Adversary Capacity

- ▶ This represents probably the most subjective quantification category. For this reason, we encourage leaning into structured (written) criteria that is defined with as much detail as team resources and bandwidth allow. Our sample criteria is provided in the next section, but we strongly encourage thoughtful adjustments according to your own understanding of and priorities within the overall threat landscape.



THREAT PROFILING PRO-TIP: CRITICAL THINKING

Analytical (critical) thinking is an essential piece of the threat profiling process. This is challenging work, involving very complex topics. If you’re in the cybersecurity field, you probably already have the chops for to meet these challenges. Don’t be discouraged by the misconception that threat profiling is only for those steeped in the threat intel discipline – you’ve got this!

- ▶ When assessing a particular threat, it is often most practical to first orient to the middle score (for example, a 3 out of 5), and then make adjustments up or down accordingly based on available evidence. As more results are scored, a comparative approach also becomes practical (e.g., threat X is relatively more capable than threat Y (currently a 3 out of 5), so threat X will receive a score of 4).
- ▶ We recognize that scoring will often be influenced by bias around information availability. For example, evidence of frequency of attacks is more a measure of likelihood or prevalence than an indication of capability. However, at times (or often), we will not have information pointing to a given adversaries' capabilities, in which case we recommend assigning a lower weight.
- ▶ For very high-level relative comparisons between nations' cyber capabilities, see the Harvard Belfer Center's National Cyber Power Index project.²⁵ Paul Jaramillo's ACTORS model offers detailed criteria that can be used to granularly measure actor "sophistication", while also calling attention to the fact that adversary sophistication may not always translate directly to effectiveness (and vice versa).²⁶

Weighting	Level	Criteria	Representative Examples
5	Superior	Characterized by groups suspected of possessing near-unlimited or very large supplies of resources. Groups often consist of many operators who generally possess high levels of skill and OPSEC. Funding is typically high and provided by a state, but may be supplemented with illicit sources. Often uses custom, sophisticated tooling (alongside existing tools) and has usually been associated with multiple novel techniques or exploits.	The most advanced/prolific APTs (e.g. APT28, Lazars Group)
4	High	Characterized by groups suspected of possessing very large resource supplies. Group members generally possess high levels of skill and OPSEC. Funding is relatively high and may be provided by a state or illicit sources. May use custom, sophisticated tooling alongside existing tools, and might be known to periodically use novel techniques or exploits.	-Major/well-known APTs supporting major adversarial nations (e.g. APT41, Fox Kitten) -The most advanced/prolific ransomware-as-a-service operations (e.g. LockBit, ALPHV/BlackCat)
3	Moderate	Characterized by possessing access to many resources, including funding which may come from a nation-state or illicit means. These groups may be linked to a considerable volume of attacks but may also have mixed levels of success and/or periodic OPSEC blunders. May use custom tooling, but it typically does not display extreme sophistication. (This is also a common assignment for APTs and major crimeware operations when knowledge gaps remain.)	-Many APTs -Many prolific initial access threats (e.g. QakBot, SocGhosh, Emotet)
1-2	Low/Limited	May be individual actors or groups, generally smaller and/or loosely organized ones. Adversaries here may claim or threaten attacks often but do not consistently follow through, at least successfully. Funding is usually limited and not at nation-state scale. Operators and their tools are usually not highly sophisticated, although some successful attacks may have occurred. Custom tools and novel exploits are uncommon. This is also a common assignment when significant knowledge gaps remain.	-Hacktivists -Lower-tier APTs & ransomware groups (including where knowledge is limited) -Infostealer campaigns

Figure 16: In an effort to be prescriptive, we have provided here our generalized perspective on weighting criteria for overall adversarial capability levels. We expect that most teams will have their own perspective on how these criteria are defined and may want to modify them and the weighting scale accordingly. For a real-world example, see the NIST SP 800-30 Rev. 1, Guide for Conducting Risk Assessments (Table D-3).³⁰

- ▶ Remember to not refrain from weighting adversaries with low scores, a tendency that we often observe in practice. It is acceptable to have adversaries in your profile that have low weighting values – this usually means simply that they do have proximity relevance but may not pose a major threat currently. These are worth monitoring for potential changes in their capabilities, which would drive the threat level higher.

Capability Density

As you likely noticed as you were building lists of identified capabilities, you will begin to see reoccurring discrete (Sub-)Techniques quite quickly after adding even a few Groups, Campaigns, or Software to your lists. As your lists grow, it is common to have considerable overlap among at least certain of the (Sub-Techniques) in your final product.

This phenomenon – which we refer to as technique “density” – is a great starting point for (Sub-) Technique quantification and prioritization. Ultimately, for defensive remediation, the most important data points generated during our profiling efforts are the discrete behaviors associated with the threats we care about, since ATT&CK allows us to directly translate from common descriptions of those techniques into relevant defensive capabilities. By focusing on the ones with the highest density, you can prioritize around the behaviors most commonly linked to your *entire* set of profiled adversaries. This in turn drives higher return on investment of perpetually limited defensive resources (budget, headcount, and technologies), while allowing you to potentially reinforce defenses against multiple discrete adversaries you care about at once.

Focus on technique density also helps us take manageable action from our threat profile – you will also notice that usually the final list of identified (Sub-)Techniques can be quite extensive. There are currently around 600 discrete Techniques and Sub-Techniques in the ATT&CK knowledge base – far fewer than the number of potential IPs from which an adversary could launch an attack or the number of vulnerabilities they could exploit, but still a large amount to try to address, especially at one time. Focus on the techniques with higher density helps deprioritize certain behaviors that might be less likely to be observed or less relevant to your specific organization, keeping you focused on committing defensive resources where they have the greatest impact.

Capability Weighting

As we noted above, you may have opportunities to surface observed victim sector or location information tied to Software, although threat intelligence sources tend to provide this metadata relatively less often than for Groups or Campaigns. However, if you do have it, you could use it for weighting in a similar manner as used for Groups and Campaigns above. Be

careful, however, to account for potential double-counting of these weights when considering a Software used by an adversary that was already weighted for sector or location factors.

Organizational context factors may have Software weighting implications, too – most notably, if your organization determined that it is especially susceptible (or concerned about the impact of) threats to a given node of the CIA Triad, that could contribute to higher weightings for certain categories of Capabilities (for example, spyware and information stealing threats will by nature pose a greater threat to the Confidentiality of data, while disruptive threats like ransomware or wipers will generally pose a greater threat to the Availability and Integrity of it, respectively). If categorizing capabilities, remember to consider that many of today's commonly used Software possess multiple capability types (some, like offensive security tools such as Cobalt Strike, are specifically crafted with modularity in mind).

Weighting around adversary behaviors is an especially important piece of the quantification process but one that often poses challenges for many teams. Many who aren't especially familiar with the ATT&CK knowledge base might not realize that certain (Sub-)Techniques will be inherently more or less notable for individual teams based on factors such as how broadly they are defined in ATT&CK and how important they are to an overall attack's execution, scope, and impact.

With so many techniques in the knowledge base, it can be intimidating to know where to begin when weighting them. For this reason, we encourage considering the overall Tactic first, and then potentially individual (Sub-)Techniques (or buckets of them) under them. We have listed several of the potential weighting factors that can be taken into account below. (Remember, this exercise just covers the threat aspect of our efforts – other important factors should be considered when evaluating defensive capability weightings as well.)

- ▶ The Tactic's or (Sub-)Technique's impact to elements of the CIA Triad if it were to be used in your environment
- ▶ The (Sub-)Technique's definition scope may have an impact on the ability to align it with discrete defensive capabilities
- ▶ Centrality to common attack chains or sequences
- ▶ Prevalence in intelligence data
- ▶ The (Sub-)Technique's influence on an attack's scope (e.g., techniques enabling an attack to spread could be weighted higher)
- ▶ The (Sub-)Technique's influence on an attack's severity (e.g., techniques that elevate privileges, which by nature could grant access to higher-value information, or techniques that impede defenders' ability to detect or respond to an ongoing incident)

PRIORITIZATION AMONG QUANTIFIED THREATS

The factors outlined above, when quantified, serve as weightings or modifiers that help drive up or down a final score, which we will use to prioritize (rank order) our list of relevant adversarial behaviors, ultimately using this to drive priority in subsequent defensive actions.

- ▶ This step should be straightforward once the preceding work has been completed. At this point, the remaining task involves simply sorting (rank ordering) from high to low the list of threats that you’ve generated thus far!
- ▶ While we have taken many steps to drive consistency and validity in our research and quantification efforts so far, we discourage rigid interpretation of the final results. This exercise is best at spotlighting significant differences between threats (for example, a “very high” versus a “medium” versus a “very low” threat), and less at distinguishing minute nuances in final scores. However...
- ▶ ...We will note that, for ease of others’ consumption and interpretation of final results, many teams will choose to divide their results list into “bands”, often associated with colors of their choice, to denote levels of overall threat. The thresholds for these bands are entirely up to your team’s decision.




Figure 17: Quantification enables evidence-based rank-ordering of threats, driving confident prioritization.

CASE STUDY: QUANTIFYING A PHARMACEUTICAL MANUFACTURER’S THREATS

Here we apply all of the research we have conducted thus far. The workflow in general involves generating priority levels for Groups, Campaigns, and Software, and using those to influence weighting of all the associated (Sub-) Techniques that we also identified.²⁷

Adversary or Campaign Name	Proximity Tier	Proximity Score	Adversary Capacity Score	Priority Tier (Includes Analyst Assessment)
Wizard Spider	Direct Threat	5	5	1
Kimsuky	Proximate Threat	4	4	1
APT29	Proximate Threat	4	5	1
Andariel	Direct Threat	5	4	2
QakBot H2 2022-H1 2023 Campaigns	Indiscriminate Threat	2	3	2
APT42	Proximate Threat	4	3	2
APT41	Proximate Threat	3	4	2
menuPass	Proximate Threat	3	4	2
Recent RedLine Stealer infection campaigns	Indiscriminate Threat	1	3	2
TA1337	Direct Threat	5	1	3
APT9	Proximate Threat	3	3	3
FIN4	Proximate Threat	3	3	3
WildNeutron	Proximate Threat	3	2	3

Figure 18: The right-hand column depicts the final results of the adversary quantification process, where the (truncated) list of pharmaceutical adversaries are rank-ordered into three priority bands after all other input factors and a final analyst assessment have been incorporated.

To begin, we assigned scores that aligned with the proximity levels from which we surfaced each Group or Software earlier (Direct, Proximate, and Indiscriminate threats). We next conducted quick assessments based on Suspected Attribution country metadata in Tidal's Community Edition and the additional details and supporting evidence provided in each adversary's profile. We combined these inputs to generate estimative priority tiers.

Close readers will notice that the QakBot and RedLine Stealer campaigns received manual bumps driven by our own analytical assessment of additional important factors that aren't currently captured with structured criteria elsewhere. Namely, since in this case we're relying on public reporting only, we don't have consistent monthly or annual infection metrics for either of these threats; however, based on the scale of the recently reported campaigns, we judged these campaigns would likely be worthy of further consideration for upcoming defensive efforts, thus manually raising them into higher priority tiers. Recency of relevant activity was also a factor, leading us to raise Kimsuky into a higher tier and lower Andariel. Individual teams may have further time-based structured criteria they choose to implement (e.g., lack of observation within a particular timeframe could automatically impact a threat's weights). These adjustments represent yet another important reminder of the importance of critical thinking when working to quantify complex topics, and of not being overly rigid with adherence to the useful but ultimately estimative structural guidelines we've established so far.

Next, we assigned similar priority levels to the Software we identified. We began by rank-ordering the list by the number of associated Groups and Software, which reveals where there may be greater overlap and, by extension, potential use of particular Software. This provided a strong overall foundation, upon which we layered another analytical assessment according to our fictional organization's priorities and concerns. Given the scale of reported activity related to LockBit 3.0 and QakBot, we raised them to higher levels, with an additional bump for LockBit considering our assessment of the serious impact to our production operations caused by a potential ransomware infection.

Finally, we move to the complete list of Techniques and Sub-Techniques associated with all of the relevant Groups, Campaigns, and Software we identified. We generated final priority levels by combining each techniques' density (how many threat entities it was associated with), the relative priority levels of those threats (generated via the previous workflows above), and finally, our fictional organization's own weightings for technique priority, which emphasized Tactics and Techniques that could increase threats to our priorities around Confidentiality and Availability of data.

Software	Associated Groups/Campaigns	Priority Tier (Includes Analyst Assessment)
Mimikatz	5	1
LockBit 3.0	1	1
QakBot	1	2
PsExec	4	2
Cobalt Strike	4	2
AdFind	3	2
PlugX	2	2
BloodHound	2	2
Empire	2	2
Net	4	3
Ping	3	3
certutil	2	3
PowerSploit	2	3
pwdump	2	3
gh0st RAT	2	3
ipconfig	2	3

Figure 19: Priority banding after the quantification assessment for the surfaced Software.

Technique	Technique ID	Tactic	Associated Groups/Campaigns/Software	Priority Tier (Includes Analyst Technique Weighting)
Winlogon Helper DLL (Boot or Logon Autostart Execution)	T1547.004	Persistence, Privilege Escalation	1	1
Multi-Factor Authentication Interception	T1111	Credential Access	1	1
Inhibit System Recovery	T1490	Impact	1	1
Domain Accounts (Valid Accounts)	T1078.002	Initial Access, Persistence, Privilege Escalation, Defense Evasion	4	2
Service Stop	T1489	Impact	2	2
Browser Extensions	T1176	Persistence	1	2
Token Impersonation/Theft (Access Token Manipulation)	T1134.001	Privilege Escalation, Defense Evasion	2	2
Registry Run Keys / Startup Folder (Boot or Logon Autostart Execution)	T1547.001	Persistence, Privilege Escalation	13	2
Windows Service (Create or Modify System Process)	T1543.003	Persistence, Privilege Escalation	12	2
Scheduled Task/Job	T1053	Execution, Persistence, Privilege Escalation	1	2
Default Accounts (Valid Accounts)	T1078.001	Initial Access, Persistence, Privilege Escalation, Defense Evasion	1	2
Dynamic Resolution	T1568	Command and Control	1	3
Proxy	T1090	Command and Control	1	3
Automated Exfiltration	T1020	Exfiltration	1	3
Group Policy Preferences (Unsecured Credentials)	T1552.006	Credential Access	1	3
Command and Scripting Interpreter	T1059	Execution	3	3

Figure 20: The final output of a successful threat profiling exercise – a list of quantified, relevant adversary techniques rank-ordered to inform defensive actions.

ACTION

The final phase of our threat profiling approach involves using the prioritized list of threats and behaviors to inform defensive improvements. Here we will review a few final factors that can influence the application of your threat profiling results.

Our profiling approach leans on the MITRE ATT&CK knowledge base, since ATT&CK provides a common language, to translate from adversary behavior into defensive capabilities. Your adversarial behavior-based profile will be most immediately applicable if your organization understands how those defensive capabilities also align with the knowledge base.

We continue to see a growing number and breadth of defensive (and offensive security)-related resources being mapped to ATT&CK. These include vendor products and open-source tools and resources, providing mitigations, protections, detections, response, logging, testing, and technical and policy controls. Tidal makes many capability mappings freely available to the community in the public Tidal Product Registry: <https://app.tidalcyber.com/vendors>

Leadership priorities and other practical considerations will also inform (and often balance against) the application of your threat profiling results. Business priorities independent of the threat landscape can and will impact focus on certain classes of security projects (whether process or technology focused), so defenders are therefore advised to consistently keep leadership-defined priorities in mind as they approach their work (and make effort to surface those priorities if they are not clear).

Figure 20 represents an ideal security validation loop, where defensive gaps against relevant threats are identified, gaps are closed through defensive reinforcements, and new (and existing) defenses are tested accurately in line with the relevant threats to ensure they are working as expected. While the threat profile output highlights which threats should independently receive a top focus, the reality is that many additional business and defensive considerations exist too which will influence the direction and timing of next steps. Most notably, your team might confidently assess that the organization’s defensive capabilities provide sufficient levels of defenses even against top threats, so further defensive focus on it and its techniques would not add much value. Conversely, there may be several potential identified gaps between the surfaced threats and current capabilities – the order and speed with which those can be addressed will realistically be influenced by many internal factors, including budget, current log sourcing, and general resource and bandwidth constraints. In practice the “widest” gap (according to quantification documented on paper) may not necessarily be the one that is addressed first or next.




Figure 21: Depiction of a representative security validation loop



CHAPTER 4:
ITERATING ON YOUR
THREAT PROFILE

UPDATING AND MAINTAINING YOUR THREAT PROFILE

We hope that the approach outlined here provides a far more practical and therefore achievable process for building threat profiles than existing resources around the subject. We acknowledge that the first few times completing these workflows may be intimidating and require some resource commitment. With that being said, realities of the modern threat landscape necessitate updates to organizations' threat profiles, and with increasing regularity.

In recent times, we have increasingly observed a phenomenon we describe as “TTP evolution” – adaptations and modifications to the tactics, techniques, and procedures used by adversaries.²⁸ Often times these shifts come as a response to positive developments in the defender landscape, such as implementation of multi-factor authentication and blocking of commonly abused macros in many cases, as well as general environmental and even world developments, such as changing (underground) and legitimate economic conditions and the war in Ukraine. While it remains that adversaries inherently don't change their behaviors as regularly as they do their attack infrastructure, the increasing pace of TTP evolution requires recognition that organizational threat profiles must also be updated to accurately account for these dynamics (frequent changes in defensive capabilities should also be reflected in capability mappings in similar fashion).

What is the appropriate update cadence? It depends; resources, bandwidth, and priorities will dictate how often teams can revert to perform threat profile updates, and the size and nature of an organization may further influence its exposure to changes in the threat landscape. At minimum, most organizations should revisit a manual threat profiling approach like the one outlined here once per year, but in most cases more frequent updates are recommended as resources permit. Twice annual and quarterly refreshes are advisable, but the increasing pace of adversary and TTP evolution point to a need for even more frequent updates wherever possible (this cadence necessitates taking advantage of automation opportunities).

An alternative approach involves a mindset and corresponding process shift towards a posture of regular threat profile maintenance (as opposed to complete refreshes at frequent intervals). This approach allows us to apply another final popular threat analysis framework, the OODA (Observe, Orient, Decide, Act) Loop. Our process for originally building a profile in many ways followed similar structure (flipping the O's), where we oriented the exercise

around our unique organizational context, then researched (observed) relevant threats, and used quantification to inform decisions on where to take action first. Threat profile maintenance involves regular evaluation of a threat's relevance *as it observed*, usually immediately after a SOC or CTI analyst collects information around some given threat activity. Using previously discussed models like the CIA Triad and Diamond Model helps to orient around organizational relevance and drive a decision to include or exclude the threat in an existing profile (update it if it is already present). Quantification weights can then all be modified accordingly, and if tools (even a spreadsheet) and/or automation are in place, final threat priority scores and rankings will update automatically.²⁹

MATURITY OPPORTUNITIES

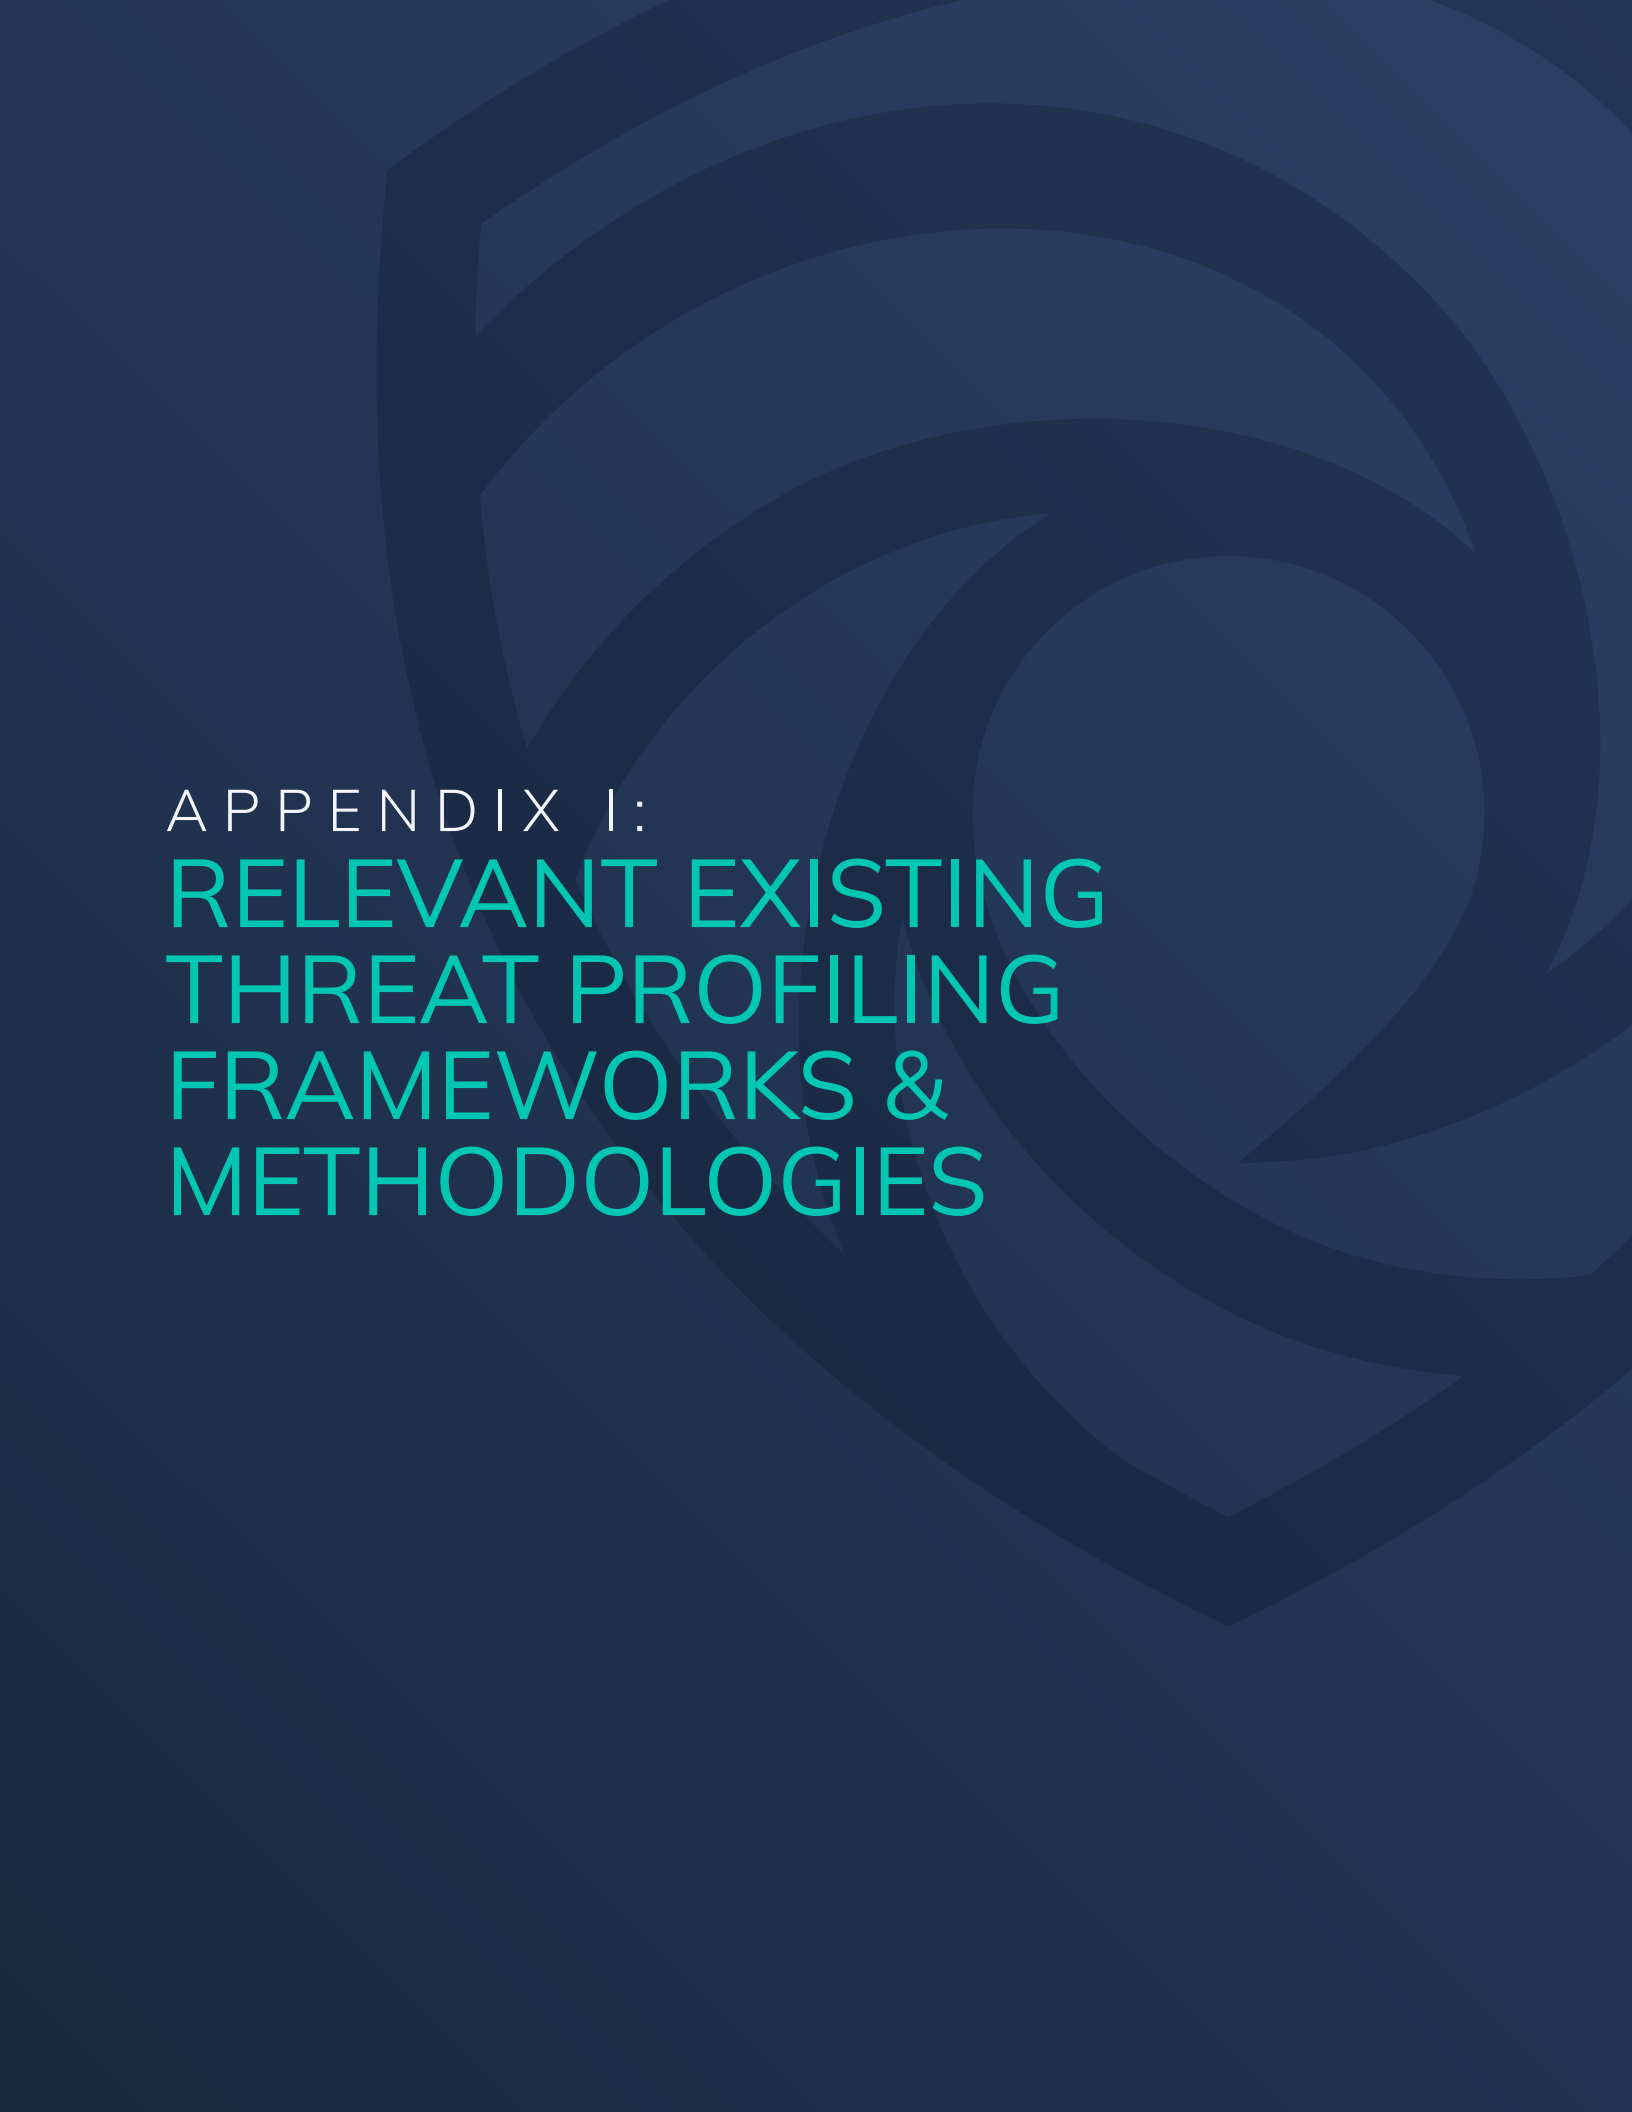
We finally want to spotlight a few opportunities for maturing your threat profiling practices, which you can consider over time as your familiarity with the approach grows and if team resources or bandwidth expand. These are derived from actual practices and workflows we have observed teams implementing in the field:

- ▶ **Measure threat profiling outputs and outcomes:** As you update (or maintain) your threat profile, we recommend tracking metrics around key outputs and results, to be able to quantitatively demonstrate change and ideally security improvements over time. Key metrics can include changes in overall and average threat priority levels (and measurements around any of the lists of threats that comprise them); a measure of how overall threat priority levels align with defensive capabilities; and specific defensive actions taken in line with or as a result of your threat profiling efforts (e.g., we wrote and/or tuned seven detection rules deployed in our EDR to address the five top techniques at the top of our profile generated in the first quarter).
- ▶ **Multiple threat profiles:** We have entirely focused on building a single threat profile for your organization as a whole, a great, practical starting point for most teams. However, additional threat profiles for certain segments of your organization or its partners can provide further granularity around the threats aligned with their unique characteristics (adding all possible threats to a single profile will likely make it difficult to manage). Popular segments that yield distinguishable differences in profiles include: geographic regions, business units/divisions, units that use unique technology stacks, and an organization's entire or segments of its third-party partners/supply chain.
- ▶ **Additional weighting factors:** As we discussed in the final quantification stage

(and addressed with manual analysis and adjustments at the time), additional factors could be introduced to add further structure and granularity around threat quantification. A few notable factors include: time-bounding analysis or “aging-out” threats and/or techniques that have been observed recently; layering both sector and location metadata (e.g. a Group is only added if it was observed targeting pharmaceutical manufacturers in the United States); consistent measurements or estimations of attack activity, prevalence, or likelihood (a persistent challenge across the CTI industry!); and financial measurement of attack impacts.

- ▶ **Automation:** Many opportunities exist to streamline elements of or entire workflows covered here. We link to many helpful public tools for interacting with ATT&CK-related datasets in the GitHub repository provided in Appendix II. Spreadsheets, scripts, and dashboarding tools may facilitate organizing, tracking, and updating lists of threats, techniques, and associated weighting scores. Advanced opportunities involve automating collection and ingestion of relevant threat intelligence (mentioning Groups, Campaigns, Software and/or Techniques) and correlating it with entities that appear within your threat profile, which supports more regular maintenance. Tidal Cyber’s Enterprise Edition can also help with a built-in threat profile builder and daily notifications of changes to the techniques being used by adversaries in your threat profile(s).





APPENDIX I:
RELEVANT EXISTING
THREAT PROFILING
FRAMEWORKS &
METHODOLOGIES

A few existing threat profiling-adjacent frameworks and methodologies, which we have observed most often in our conversations with security practitioners, are listed below for reference. A 2018 review of key existing frameworks and methodologies for “threat modeling”, published by the Homeland Security Systems Engineering and Development Institute (HSSEDI)[™] (operated by The MITRE Corporation), also contains a roundup that includes many of these and several additional resources: <https://www.mitre.org/sites/default/files/2021-11/prs-18-1174-ngci-cyber-threat-modeling.pdf>

- ▶ Enterprise Threat Model Technical Report: <https://www.mitre.org/sites/default/files/2021-11/pr-18-1613-ngci-enterprise-threat-model-technical-report.pdf>
- ▶ Process for Attack Simulation and Threat Analysis (PASTA): <https://versprite.com/blog/what-is-pasta-threat-modeling/>
- ▶ Guide for Conducting Risk Assessments: <https://nvlpubs.nist.gov/nistpubs/legacy/sp/nistspecialpublication800-30r1.pdf>
- ▶ STRIDE: [https://learn.microsoft.com/en-us/previous-versions/commerce-server/ee823878\(v=cs.20\)](https://learn.microsoft.com/en-us/previous-versions/commerce-server/ee823878(v=cs.20))
- ▶ DREAD: <https://adam.shostack.org/modsec08/Shostack-ModSec08-Experiences-Threat-Modeling-At-Microsoft.pdf>
- ▶ LINDDUN: <https://people.cs.kuleuven.be/~kim.wuyts/LINDDUN/LINDDUN.pdf>
- ▶ Factor Analysis of Risk Information (FAIR[™]): <https://www.fairinstitute.org/what-is-fair>
- ▶ Trike: <http://www.octotrike.org/>
- ▶ Visual, Agile and Simple Threat (VAST): <https://threatmodeler.com/threat-modeling-methodologies-vast/>
- ▶ Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE[®]): <https://resources.sei.cmu.edu/library/asset-view.cfm?assetID=8419>



APPENDIX II:
THREAT PROFILING
RESOURCES

A list of excellent general threat profiling resources – which we have consistently referenced over years of developing the guidance provided here – is shared below.

A much larger library of more tactical resources and helpful reference materials is hosted in a Tidal GitHub repository purpose-built to accompany this paper: <https://github.com/tidalcyber/cyber-threat-profiling>

- ▶ Using Threat Intelligence to Focus ATT&CK Activities: <https://www.youtube.com/watch?v=V--wxuSEMDO>
- ▶ How to prioritize effectively with Threat Modeling and ATT&CK: <https://www.youtube.com/watch?v=i5mx8jyoOGE>
- ▶ Resistance Isn't Futile: <https://www.youtube.com/watch?v=b0ShMaKDidU>
- ▶ Hunting for Post-Exploitation Stage Attacks: <https://www.youtube.com/watch?v=PdCQChYrxXg>
- ▶ Adversarial Threat Modelling: https://github.com/ssnkhan/adversarial-threat-modelling/blob/master/Adversarial-Threat-Modelling_Presentation.pdf
- ▶ Quantifying Threat Actors with Threat Box: <https://klgrz.medium.com/quantifying-threat-actors-with-threat-box-e6b641109b11>
- ▶ Sophisticuffs: The Rumble Over Adversary Sophistication: <https://www.slideshare.net/PalJaramillo/bsides-chicago2017>
- ▶ Getting Started with ATT&CK: Threat Intelligence: <https://medium.com/mitre-attack/getting-started-with-attack-cti-4eb205be4b2f>
- ▶ Using ATT&CK for CTI Training: <http://attack.mitre.org/resources/training/cti/>
- ▶ Emulation Planning for Purple Teams: <https://www.academy.attackiq.com/courses/emulation-planning-for-purple-teams>



CYBER THREAT PROFILING GLOSSARY

For consistency throughout this guide, definitions for several key terms relevant to the threat profiling discipline are provided below. We acknowledge that different corners of the community continue to use often overlapping (and sometimes competing) versions of these definitions and that readers might choose to use their own variations – and that’s great! We encourage thoughtful consideration of these complex topics. We believe the most important point is consistency in your own team’s use of these terms (ideally you have documented your internal definitions too).

Threat-Informed Defense: The systematic application and deep understanding of adversary tradecraft and technology to assess, organize and optimize your defenses.

Cyber Threat Profiling: A structured, repeatable process for determining relevant, prioritized cyber threats (adversaries, malware, & associated attack techniques), based on quantifiable evidence.

MITRE ATT&CK®: According to its website, MITRE ATT&CK® (ATT&CK) “is a globally-accessible knowledge base of adversary tactics and techniques based on real-world observations”.⁴ ATT&CK stands for “Adversarial Tactics, Techniques and Common Knowledge”.

Enterprise: In this context, **Enterprise** refers to virtually any organization, public or private (it generally denotes a relatively large organization). Tidal’s approach to threat profiling is distinct from many existing threat profiling/modeling methodologies because it focuses primarily on surfacing threats to organizations as a whole, as opposed to individual assets or systems (groupings of assets).




Figure 22: A visual representation of the key components comprising the concepts of “Threat” and “Risk”, some of the many terms often confused when discussing cyber threat profiling.

Adversary: Typically refers to an attacker **Group**, a defined cluster of related cyber threat activity. Often used interchangeably with the terms **Actor** or **Threat Actor**.

Campaign: Refers to a set of cyber threat activity observed in the real world, which takes place within a given period of time. Campaigns are carried out by identified or unknown **Groups**.

Software: Any computer code or program. In **ATT&CK** lexicon, Software is divided into two types: **Tool** (a legitimate, benign utility, often abused by adversaries for malicious purposes) and **Malware** (software specifically crafted for malicious purposes).

Tactic: The reason or objective (the “ends” or the “why”) behind an attacker action.

Technique: The means by which (the “how”) an attacker action is carried out.

Sub-Technique: In **ATT&CK** lexicon, a more specific description of a **Technique**.

Procedure: The specific implementation of an **ATT&CK Technique** or **Sub-Technique**.

TTP: A collective term referring to activity that comprises a **Tactic**, **Technique**, and **Procedure** (“TTP”).

Behavior: A collective term referring to activity that comprises a TTP and involving one or more **Platform(s)**.

Platform: Technology categories to which a technique is applicable

CTI: Cyber Threat Intelligence

CIA Triad: A popular approach for categorizing the foundational components

of information security risks, comprising the data properties of **Confidentiality**, **Integrity**, and **Availability**.

Confidentiality: In the context of the CIA Triad, a property of data where they are only accessible by the individual(s) to which the data’s owner intends to grant access.

Integrity: In the context of the CIA Triad, a property of data where retrieved data arrives in its original state.

Availability: In the context of the CIA Triad, a property of data where they can be accessed when and how the data’s owner intends.

Diamond Model (of Intrusion Analysis)⁵: According to its foundational whitepaper, published by Sergio Caltagirone, Andrew Pendergast, and Christopher Betz in 2013, the Diamond Model “establishes the basic atomic element of any intrusion activity, the event, composed of four core features: adversary, infrastructure, capability, and victim”. Each node in the model is “edge-connected” which represents underlying relationships among the four features.

Threat: A hazard that could cause harm. The classic, academic, mathematical representation of Threat is $\text{Threat} = \text{Intent} \times \text{Capability} \times \text{Opportunity}$

Motivation: A collective representation of an adversary’s objective(s).

Intent: An adversary’s desire to attack a potential victim(s).

Capability: Adversarial Capability refers to the collective means an attacker possesses to harm potential victims. **Software** and **Techniques** represent specific potential

Adversarial Capabilities. A **Defensive Capability** refers to an asserted and/or validated ability to defend against a specific technique.

Opportunity: A threat **Opportunity** is realized when there is alignment between an adversary's **Motivation** and victim(s)' characteristics. Time, space, and **Capability** factors must also align.

Threat (or Risk) Register: A list of discrete threats (or risks) that an organization deems relevant.

Attack Surface: The collective physical and technological footprint of organizational assets that adversaries could attack.

Vulnerability: In the context of threat profiling, **Vulnerability** represents a conceptual condition where an asset's or organization's relevant **Defensive Capabilities** do not fully protect its **Attack Surface**.

Quantitative: Can be measured in specific, exact, and defined terms. Contrasts with **Qualitative**, which refers to a subjective or estimative approach.

Magnitude: The size or scale of something.

Likelihood: The mathematical/quantitative chance that an event will take place. Also known as **Probability**.

Impact: The collective consequences of an event, typically defined in quantitative and especially financial terms.

Risk: Generally represents a combination (or balancing out) of a **Threat** assessment and compensating factors. The mathematical

representation of Risk used here is **Risk = (Threat) Likelihood x Impact**.

Observed (Threat): An instance of a **Threat** that has been witnessed (and usually documented), publicly or privately.

Targeted (Threat): An instance of a **Threat** that possesses defined (usually assessed) **Intent**.

Threat Modeling: In most practical ways, we find that this term is largely synonymous with Threat Profiling, although it usually carries a more mathematical connotation and is most regularly associated with threat assessments involving individual assets (especially web applications).



ENDNOTES

- 1 *Tidal defines threat-informed defense as: "The systematic application and deep understanding of adversary tradecraft and technology to assess, organize and optimize your defenses."* <https://www.tidalcyber.com/blog/threat-informed-defense-what-is-it>
- 2 <https://mitre-engenuity.org/cybersecurity/center-for-threat-informed-defense/>
- 3 <https://mitre-engenuity.org/cybersecurity/attack-evaluations/>
- 4 <https://attack.mitre.org/>
- 5 <https://www.activeresponse.org/wp-content/uploads/2013/07/diamond.pdf>
- 6 Inability to rank-order threats ultimately means that all threats on your list are priorities, which in our experience usually means none of them are.
- 7 <https://www.tidalcyber.com/blog/adversary-ttp-evolution-and-the-value-of-ttp-intelligence>
- 8 <https://www.securitymagazine.com/articles/98776-one-of-the-biggest-threats-to-a-cybersecurity-team-employee-burnout>
- 9 <https://medium.com/@girishsj2/threat-modeling-the-buzz-word-5c8c9d475e0f>
- 10 Readers will notice we refer to a threat profiling "approach" and not a framework or methodology – this is intentional. Data points surfaced at various phases will inform work at other stages, and teams may choose to begin at different entry points. The complexity of existing methodologies has limited wider profiling adoption, so we encourage tailoring the workflow to the point it can be completed given your unique resource and experience levels.
- 11 As a reminder from the Glossary, we use the term "Enterprise" to distinguish our profiling approach from others that focus on assets or collections of assets (systems). Despite connotations that might suggest the private sector, our approach is absolutely applicable to threat profiling around public/government sector agencies/entities. (Note that our approach was not necessarily scoped for country-level assessments, although we'd love to hear if you have success with such an application.)
- 12 <https://www.activeresponse.org/wp-content/uploads/2013/07/diamond.pdf>
- 13 We also acknowledge that managed security providers increasingly provide profiling services to their clients, so the guidance in this section is especially relevant in cases where the individual completing the exercise may not even be employed by the organization they are assessing.
- 14 SEC filings can be searched here: <https://www.sec.gov/edgar/searchedgar/companysearch>
- 15 TeamTNT Group profile: <https://app.tidalcyber.com/groups/325c11be-e1ee-47db-afa6-44ac5d16f0e7-TeamTNT>
Resource Hijacking Technique profile: <https://app.tidalcyber.com/technique/d10c4a15-aeaa-4630-a7a3-3373c89a584f-Resource%20Hijacking>
- 16 <https://www.cfr.gov.au/publications/policy-statements-and-other-reports/2020/corie-pilot-program-guideline/pdf/corie-framework-guideline.pdf>
- 17 <https://www.youtube.com/watch?v=i5mx8jyoOGE>
- 18 A note on Opportunity: According to our research, most academic definitions of "threat" include a measure of Opportunity. In our experience, this is one of the least practical for organizations to assess, since it involves a difficult-to-measure estimation of time and space alignment of adversary intent plus the existence of a relevant vulnerability. We have seen teams use a truncated definition of Threat = Intent x Capability effectively (and conversely, have seen little practical guidance for measuring true Opportunity), and so we chose to exclude it from the guidance provided here.
- 19 Teams may choose to skip this step – and that's ok. In addition to the attribution challenges detailed above, we have observed many teams that intentionally seek an independent assessment of their profile based entirely on externally sourced data, which provides a check against potential bias.
- 20 <https://cybersecurity.att.com/blogs/security-essentials/understanding-malware-as-a-service-maas-the-future-of-cyber-attack-accessibility>
- 21 We find that this terminology best captures the phenomenon described here. A close and often overlapping term is "opportunistic", but there are important academic differences between the two.

- 22 <https://www.tidalcyber.com/blog/ransomware-threat-profiling-prioritizing-indiscriminate-threats>
- 23 <https://klgrz.medium.com/quantifying-threat-actors-with-threat-box-e6b641109b11>
- 24 <https://nvlpubs.nist.gov/nistpubs/legacy/sp/nistspecialpublication800-30r1.pdf>
- 25 <https://www.belfercenter.org/publication/national-cyber-power-index-2022>
- 26 <https://www.slideshare.net/PaJaramillo/bsides-chicago2017>
- 27 Note that for space considerations, we aren't showing the full results of any of the Group/Campaign, Software, or Technique lists here. The limited scope covered here also led us to use a three-point scale for the final "priority tiers", but we recommend considering a five-point scale that provides appropriate granularity for many profiles.
- 28 <https://www.tidalcyber.com/blog/adversary-ttp-evolution-and-the-value-of-ttp-intelligence>
- 29 We launched the Making Waves blog series to help further demonstrate the pace of (publicly observed) TTP change and support defenders in identifying where there might be notable shifts or trends in the landscape that could inform their threat profiles - down to the individual adversarial technique level: <https://www.tidalcyber.com/blog/making-waves-ttp-intelligence-highlights-in-march>
- 30 <https://nvlpubs.nist.gov/nistpubs/legacy/sp/nistspecialpublication800-30r1.pdf>



ABOUT TIDAL CYBER



Tidal Cyber makes threat-informed defense achievable for organizations of all sizes. The Tidal Platform helps our customers map the security capabilities of their unique environment against the industry's most complete knowledgebase of adversary tactics and techniques including the MITRE ATT&CK® knowledge base, additional open-source threat intelligence sources, and a Tidal-curated registry of security product capabilities mapped to specific adversary techniques. The result is actionable insight to track and improve their defensive coverage, gaps, and overlaps.

COMMUNITY EDITION



Tidal's Community Edition is the freely-available threat-informed defense platform for researching threat actors, building technique sets, and so much more. Community Edition Users are able to share their work and participate in the larger Tidal Cyber community of defenders.

ENTERPRISE EDITION



Tidal Enterprise Edition brings a full-featured threat-informed defense experience to large enterprises and security teams. By pairing the threats most relevant to the organization with the tools in an organization's defensive stack, Tidal Enterprise Edition gives a complete picture of an enterprise's cyber posture, and quantifies how confident the organization can be in the Tidal Confidence Score™.

ABOUT THE AUTHOR



Scott Small is Tidal Cyber's Director of Cyber Threat Intelligence. He has spent the large majority of his career focused on threat quantification, originally assessing risks for physical security issues like terrorism & drug trafficking and now focused on cyber threats. During his time supporting a large number and variety of organizations, especially enterprises, with threat intelligence analysis, he regularly witnessed struggles with the application of that intelligence. These observations led him to embrace the concept of threat-informed defense as a practical yet powerful approach for improving security efforts. He is excited (and encouraged) to witness the growing adoption of threat-informed defense approaches and mindsets, and he hopes this resource can further support that welcome trend.