



# THE THREAT HUNTER'S HYPOTHESIS

**A case for structured threat  
hunting and how to make it work  
in the real world.**



# TABLE OF CONTENTS

Everybody Wanna Threat Hunting	3
The Threat Hunter's Conundrum	6
Limitations of Unstructured Hunting	9
The Case for Structured Hunting	11
Building the Hunt	13
Why Hasn't Anyone Solved This Problem?!	16
The Threat Hunting Content Platform	20
A Threat Hunting Content Platform in Action	25
A Buyer's Guide to Content Platforms	27

# “EVERYBODY WANNA THREAT HUNT...”

There’s a saying in the bodybuilding world. It’s originally a quote from 8X Mr. Olympia, Ronnie Coleman:

**“EVERYBODY WANNA BE A BODYBUILDER, BUT DON’T NOBODY WANNA LIFT NO HEAVY-ASS WEIGHT.”**

Believe it or not, there’s a corollary to threat hunting. Every large organization wants to threat hunt, but they don’t want to pay upwards of \$150,000 per head to get it done. Many organizations try to solve their security problems by investing in technology or automation but remain hesitant to invest those same dollars in skilled threat hunters.

As a result, threat hunting teams usually lack the manpower to conduct a high volume of hunts. They are also typically left to ‘fend for themselves,’ with few (if any) tools or services to support their needs.

In this paper, we’ll look at the challenges facing threat hunting teams and why existing tools don’t solve them. We’ll make the case for ‘structured’ threat hunting as the most effective way for hunting teams to operate and explain how to make it work consistently and effectively in the real world.

## WHY THREAT HUNT?

If you *aren't* a threat hunter, you might wonder, *“why do we need to pay for threat hunting when we already spend a fortune on automated tools?”*

The answer is simple: standard ‘protect and detect’ tools don’t completely defend your organization from cyber threats. You know all those big organizations that were breached in the last year? *They had the same tools in place*—yet, they were breached. And when they finally found out, it was a *human* that uncovered the evidence.

Standard cybersecurity tools are *reactive*. Detection and protection tools are based on rules, and those rules can only be created after a threat has been observed in the wild. If you’re at a large organization—for instance, in financial services or healthcare—you could easily be among the first to be hit with a new threat. In that case, your tools provide no protection.

The 2020 Threat Hunting Report by Cybersecurity Insiders found that automated tools miss an estimated 30% of all threats—and 56% of SOCs identify detection of advanced and emerging threats as a top challenge.

Threat hunting is designed to solve these problems. It’s a proactive approach to security that searches for unknown threats in your environment. Done well, it dramatically improves your chances of spotting and removing threats before it’s too late.

## STANDING ON THE SHOULDERS OF SQRRLS

Throughout this paper, we'll refer to some of the founding contributors to the field of threat hunting. One such forerunner is SQRRL, a former platform vendor with a focus on threat hunting education. Some of their content is still available here: <https://www.threathunting.net/sqrrl-archive>

### TL;DR: SUMMARY OF THE THREAT HUNTER'S HYPOTHESIS

- ▶ Organizations increasingly recognize the value of threat hunting, but hunting teams lack personnel, resources, and tools.
- ▶ Unstructured threat hunting is a step up from reactive measures but lacks the consistency to disrupt threat groups' operations.
- ▶ Structured threat hunting is far more effective for detecting new threats—but it takes a long time to develop structured hunts.
- ▶ Current sources of threat hunting content are outdated, lack context, and require arduous customization and validation.
- ▶ A threat hunting package addresses challenges with hunting content by providing everything needed to run a successful hunt.
- ▶ A threat hunting content platform provides current, highly effective hunting packages that return actionable results and fill critical gaps in existing cybersecurity programs.

# THE THREAT HUNTER'S CONUNDRUM

The last few years have seen an explosion of interest in threat hunting. Security-conscious organizations from many industries have set resources aside to build a hunting capability.

However, most threat hunting teams face six huge challenges:

---

## CHALLENGE #1: BUDGET CONSTRAINTS

Threat hunting is often poorly defined and understood—particularly by senior leaders. It's often unclear how it fits into compliance and best practice frameworks. As a result, threat hunting is often 'last in line' for budget, receiving only a fraction of the resources given to blue and red teams. Most of this budget is spent to maximize headcount, but most hunting teams are still extremely low on skilled FTEs.

---

## CHALLENGE #2: THE SKILLS GAP

The skills gap affects all areas of cybersecurity, but few as keenly as threat hunting. There are few people worldwide who call themselves threat hunters, and even less with a full skillset. As a result, most teams are low in headcount and can't run nearly as many hunts as they would like.

---

## CHALLENGE #3: LACK OF DEDICATED RESOURCES

The SANS 2020 Threat Hunting Survey found 75% of organizations with a threat hunting capability use staff who have other roles—usually in a SOC and IR team. This is preferable to no threat hunting but makes it hard to define consistent, repeatable processes and build effective hunts.

Many SOC and IR analysts are fully capable of developing threat hunting skills. However, threat hunting is a different profession with its own certifications and qualifications. When forced to split their time between both functions, there is very little remaining for analysts to develop their skills and obtain the necessary certifications.

Worse still, most SOC and IT teams are already challenged with unmanageable workloads. When forced to split their time, the organization's security posture is compromised, as issues **will** be missed.

---

## CHALLENGE #4: THREAT INTELLIGENCE IS NOT GEARED TO THREAT HUNTING

Most organizations with a threat hunting capability have a Cyber Threat Intelligence (CTI) team in place. However, unless the CTI team is fully mature, it won't include personnel with the broad skill sets needed to produce behavioral and TTP-related intelligence to inform threat hunts.

For instance, a core capability needed to inform threat hunting is malware reverse engineering. A full understanding of what malware does and how it works is essential to develop a hunt for the TTPs it uses. However, this is an extremely uncommon skill set that most CTI teams can't retain.

Sandboxes are a common attempt to solve this problem, but lots of malware can evade this approach—leaving threat hunters without adequate intelligence to inform their hunts.

### DON'T GET IT TWISTED

Some organizations wrongly conflate CTI with threat hunting. CTI informs on threats that pose a high risk to the organization. Threat hunters search for evidence of malicious activity within the organization's environment. CTI guides threat hunting—without intelligence, a threat hunter could search forever and find nothing.

---

## CHALLENGE #5: THE 'LEGITIMACY GAP'

As IT environments grow more complex, so do behaviors observed in them. This causes a growing knowledge gap of what is legitimate, even among skilled threat hunters. A hunter might observe behavior that appears unusual, but due to changing architecture, it could be legitimate. This is a challenge for threat hunting teams, which must continually keep abreast of evolving internal conditions and behaviors.

---

## CHALLENGE #6: CONFUSED TERMINOLOGY

Across the industry, there's a lack of consistency in the meaning given to 'threat hunting.' For clarity's sake, here are the three activities most commonly described as threat hunting:

1. IoC hunting.
2. Unstructured threat hunting.
3. Structured threat hunting.

We'll give clear definitions of unstructured and structured threat hunting later in this paper. For now, we want to clear up a misunderstanding.

There is discussion on whether hunting originating from an IoC qualifies as threat hunting. It doesn't.

**By definition, threat hunting searches for *unknown* threats.** Threat hunters search for previously undetected activity tied to malicious artifacts and behaviors that cannot be found by detection capabilities or alerts.

IoCs relate to *known* threats. If a threat is known, it falls under the parameters of detection and alerting capabilities—not threat hunting.



# LIMITATIONS OF UNSTRUCTURED HUNTING

Threat hunting is the use of advanced techniques to find ‘bad stuff’ in an environment. There are two disciplines—*structured* and *unstructured* hunting. We’ll explain structured hunting in a moment.

Unstructured—also known as *data-driven hunting*—is the most common form of threat hunting. Almost half (45%) of respondents to the SANS 2020 Threat Hunting Survey said their organization takes an ad hoc approach to threat hunting.

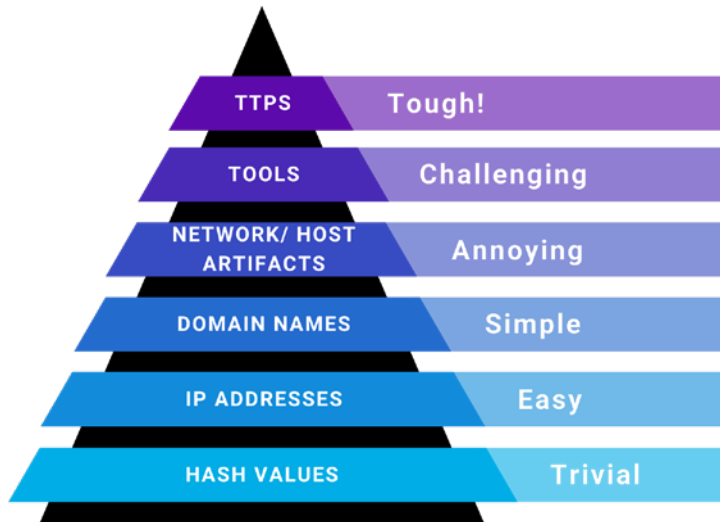
In an unstructured hunt, threat hunters dig through log files manually and opportunistically, using data manipulation techniques to find anomalies. These hunts are perfectly valid and more proactive than traditional protection and alerting.

However, due to their *ad hoc* nature, unstructured hunts aren’t repeatable or reliable. They may find valid results, but there’s no way to repurpose successful hunts in future detections. They are also unlikely to spot sophisticated threats or groups.

## CLIMBING THE 'PYRAMID OF PAIN'

The Pyramid of Pain was developed by David Bianco, former Security Technologist at SQRRL, to explain the relationship between:

1. The categories of indicators a threat hunter can use to detect malicious activity; and,
2. How much 'pain' it causes threat groups when an organization prevents them from using those indicators.



Source: The Pyramid of Pain, David Bianco, 2017.

**Note:** TTPs stands for Tactics, Techniques, and Procedures—the behaviors threat groups use to achieve malicious objectives.

Most organizations are only capable of ingesting CTI that focuses on the bottom half of the pyramid by continually adding indicators to detection tools. As these indicators are 'denied,' threat groups simply refresh their infrastructure and carry on—it costs them almost nothing.

Threat hunting aims to identify **artifacts** and **behaviors** from the **top** half of the pyramid. These are much harder for threat groups to alter. They require significant changes to infrastructure and tooling—and potentially learning completely new techniques.

Why is this relevant? Even successful unstructured hunts are only likely to identify artifacts from lower down the pyramid of pain. This is temporarily useful but doesn't provide a long-term security benefit.

Take the Ryuk ransomware. Since it was first spotted in 2018, Ryuk has gone through thousands of changes to low-level artifacts. Finding one of these in your environment is temporarily useful, but it won't protect against future adaptations.

# THE CASE FOR STRUCTURED HUNTING

If unstructured hunting mainly finds low-level artifacts, how do you find more 'painful' TTPs? That's where structured threat hunting comes in.

A structured threat hunt is a formal search for TTPs—specifically, those used by threat groups deemed to be a significant risk to the organization. It's a search for behaviors rather than indicators or artifacts.

## EVEN APTS ARE LAZY

TTPs are right at the top of the Pyramid of Pain. Even Advanced Persistent Threat (APT) groups are prone to reusing TTPs across their operations. Attackers develop these capabilities over years and may also be constrained by standard operating procedures. As a consequence, TTPs are far harder to change than infrastructure.

Returning to Ryuk, the ransomware has been associated with thousands of artifacts. Trying to uncover new Ryuk activity using artifacts from last month won't work. However, all versions of Ryuk use the same common persistence mechanism in Windows to stay active on infected systems. A structured, hypothesis-driven hunt can identify this behavior **even if the malware isn't associated with any known indicators or artifacts.**

When you deny attackers the use of their favorite TTPs, you force them to develop new behaviors and tools. This costs time and resources, and may even prompt them to look for easier targets.

## THE THREAT HUNTER'S HYPOTHESIS

Identifying malicious behaviors requires a different skillset to unstructured hunting. The structured hunting process applies custom queries to security datasets to identify possible malicious behavior. Before that, though, there's a critical preliminary step.

Each structured hunt is built around a central hypothesis about specific attackers and their associated TTPs. This hypothesis takes the form of a clear, falsifiable statement, often driven by external CTI. For example:

*Our CTI has identified that APT28 is targeting our industry. Once they establish a beachhead in an environment they begin to carry out internal reconnaissance campaigns. APT28 is known to use an automated custom script with known behaviors to conduct host enumeration as part of internal reconnaissance.*

Notice the definite language. This hypothesis is either correct or incorrect; there's no middle ground. The threat hunter researches the threat to understand its characteristics, gathers CTI, and writes a query to find specific host enumeration techniques and procedures in the environment.

A hypothesis can be derived from several different sources, including:

- ▶ CTI — e.g., prompted by information from an open or privileged source such as media coverage of an emerging threat or from government agencies.
- ▶ TTPs — e.g., to provide coverage of a tactic or technique in the MITRE ATT&CK or other frameworks.
- ▶ Risk assessment — e.g., to provide coverage of an identified 'blind spot' in the organization's security controls.

Once the hypothesis is set, the threat hunter develops a query to search for associated threat behaviors in an existing data set—for example, in the organization's SIEM, EDR, data lake, or IDS. Depending on the threat, the hunter may run this query sporadically, or they may set it to run continuously in the background.

# BUILDING THE HUNT

So, what's the catch? There are three.

First, threat hunting is a highly skilled role. A very small subset of security professionals have the expertise to research and develop threat hunting hypotheses and queries (often called hunting 'content').

Second, building structured hunts requires a top-tier CTI capability to provide the context and behavioral intelligence to fuel hypotheses. Most CTI teams aren't geared towards guiding threat hunts—this is an additional capability to develop.

Third, developing a threat hunt from end-to-end takes a long time—often two or three weeks. In practice, most threat hunting teams only have the capacity to run a handful of hunts each month.

# THE STRUCTURED THREAT HUNTING PROCESS

Why does building a hunt take so long? Because the process is complex, requiring many steps:

1. Develop an initial hypothesis based on CTI, risk, etc., and refine it into a definite, falsifiable statement.
2. Research the threat or group the hypothesis is based on to understand its characteristics and behaviors.
3. Write a query to harvest results from a security dataset.
4. Refine the query until it runs quickly and returns a manageable number of results with few false positives.

**Note:** Usually, the first version of a query is too broad, returning lots of results with many false positives. The threat hunter must then analyze the results to better understand what they are looking for.

By refining the query—for example, by identifying common traits of false positives and excluding them from the search—the threat hunter can gradually pare down the results to a manageable dataset with few false positives. This optimization process requires a deep understanding of the threat and its associated TTPs.

Query runtime is also an important consideration, and threat hunters often refine queries to run more quickly. Slow queries don't just inconvenience the hunter—they slow down the security tool for all users.

5. Test and validate the query to ensure it works as intended against the target TTP.
6. If it doesn't or only works some of the time, the hunter must identify this as it creates additional risk.
7. Produce documentation on what the query does, how it works, and how to deploy it. Documentation must be clear and comprehensive—poor or missing documentation seriously harms a hunt's effectiveness and repeatability.
8. Produce runbooks to help analysts understand and validate results from the query. Without this, analysts are forced to process results based on their personal experience and knowledge, leading to possible inconsistent results and actions.

9. Produce remediation guidance with recommended next steps. This helps analysts ensure repeatable best practice remediation whenever a threat is detected.
10. Deploy the query into a production environment, e.g., SIEM, EDR, IDS, or data lake (to name a few!).
11. Continually revise the query to keep up with changing behaviors and detect the latest version of the targeted threat. Threat actors have an incentive to remain undetected, so this is an ongoing process.

This process is laborious and time-consuming, even for a highly skilled threat hunter.

So, we're in a bind. Structured threat hunting is the best way to identify sophisticated threats and disrupt the Cyber Kill Chain... but it takes a tremendous amount of skill, energy, and resources to do successfully.

Which begs the question...

# WHY HASN'T ANYONE SOLVED THIS PROBLEM?!

If you're in the threat hunting game, you know there have been several attempts to solve these challenges. Notably:

1. Open source threat hunting content.
2. Vendor-provided threat hunting content.
3. Developing threat hunting content in-house.

The SANS 2020 Threat Hunting Survey found only around 3.4% of organizations outsource threat hunting to a third party. That means, aside from developing content in-house—which runs into the resourcing issues discussed earlier—threat hunting teams have just two sources of threat hunting content: open sources and security vendors.

## WHAT IS THREAT HUNTING CONTENT?

There is some confusion on this, so let's take a moment to clear it up.

The accepted definition of threat hunting content is a machine-readable *query*, with no context, verification, or guidance. Threat hunters can't use these queries without manually verifying and altering them to fit their environment. There's value to be had from threat hunting content as it currently stands, but there's a lot of scope for improvement.

In the next section, we'll introduce what we consider to be a vastly superior resource: the threat hunting *package*.



## OPTION #1: OPEN SOURCE CONTENT

Open source content comes from the global cybersecurity community. Threat hunting experts and enthusiasts share their expertise by making queries freely available via websites and platforms.

In our experience, open sources are the most common source of threat hunting content. However, free content comes at a price.

### OPEN SOURCE CONTENT IS...

- ▶ Not updated on an ongoing basis.
- ▶ Not vetted and may not be effective.
- ▶ Hard to validate, forcing threat hunters to expend precious time determining whether the content is worthwhile.
- ▶ Missing context and surrounding CTI.
- ▶ Lacking transparency—threat hunters have little idea how open source content was developed, making it hard to rely on.
- ▶ Not customized to each organization's environment—threat hunters must adapt content to fit their needs.

While open source content can save threat hunting teams some time, it comes with significant drawbacks. The lack of transparency and vetting means professional threat hunting teams must spend a lot of time on research, validation, and customization—perhaps less than developing content from scratch, but more than they would like.

## OPTION #2: VENDOR-PROVIDED CONTENT

Cybersecurity tool vendors are a common source of threat hunting content. Vendors of SIEM, EDR, and other security tools provide free content within their solutions, which typically takes the form of pure queries with little (if any) surrounding context, supporting intelligence, or human-readable guidance.

In addition, vendor-provided content is often...

- ▶ 'Stale' or out-of-date, making it practically useless.
- ▶ Low in sophistication and not tailored to an individual hunter's environment or technology stack.
- ▶ Lacking transparency, context, and relevant CTI.

## UNIVERSAL CONTENT CHALLENGES

Open source and vendor-provided content provide some benefits to threat hunting teams. However, they also present challenges. Unfortunately, we haven't covered all of them yet.

Fundamentally, almost all content currently available to threat hunting teams suffers from four major issues:

- 1. Queries are too 'wide' and return too many results.** They take too long to run, introduce performance issues, and force threat hunters to go through cycles of refining and rerunning queries.
- 2. Content is often out-of-date** because vendors and open sources don't use decay modeling to update or remove stale entries. This also forces threat hunters to search through an unmanageable volume of content and doesn't equip them to hunt for emerging threats.

**Decay modeling** is a mathematical formula that determines the diminishing value of a threat hunting package over time. Once a package falls below a set threshold, it should be 'decayed' out. If the package becomes relevant again in the future, it should be reintroduced.

- 3. Content is not customized to the hunter's environment.** Each organization has its own unique tool stack and configuration, which generic content doesn't accommodate.
- 4. Content is not mapped to security frameworks.** Mature hunting teams use models like MITRE ATT&CK to identify gaps and prompt threat hunts. Generic content is unmapped, making it hard to fit into a cybersecurity program—and even harder to evaluate.

## THE 'FALSE SENSE OF SECURITY' TRAP

The most dangerous aspect of low-quality threat hunting content is the false sense of security it gives organizations.

Outdated or poorly written content can lead an organization to *believe* it is protected against a threat—but the reality could be the opposite. This leaves them vulnerable to being blindsided by threats they thought were under control.

# THE THREAT HUNTING CONTENT PLATFORM

Threat hunting teams need an alternative to the simplistic, outdated content provided by open sources and tool vendors—*and* an alternative to hiring additional FTEs or ‘borrowing’ them from other security functions.

However, content is still king, whether developed in-house or by an expert external source. Without content that includes context, guidance, queries, emulation, and remediation guidance, there is no structured hunting.

Consider the challenges threat hunters face:

- ▶ They have the skill—but not the time—to develop effective, repeatable hunts.
- ▶ Using pre-written content saves a lot of time, so long as it’s current, accurate, complete, and easy to validate.
- ▶ Content provided by open sources and security tool vendors does not fit these requirements.

The solution is obvious: **Threat hunters need access to a library of threat hunting content that is current, accurate, complete, and easy to validate.** That’s where threat hunting content platforms come in.

## WHAT IS A THREAT HUNTING CONTENT PLATFORM?

A content platform gives threat hunting teams access to a complete library of content—continuously updated and vetted—covering a wide range of current threats.

Delivered via a web-based portal, hunters get access to fully built-out threat hunting content developed by 'best of the best' threat hunters. Fully tagged and easily searchable, hunters can quickly identify content that supports objectives and fills gaps in their cybersecurity program.

A threat hunting content platform provides content tailored to the customer's environment. No more time wasted on adjusting queries to different tools or configurations. It's all done automatically by the platform.

Most hunting teams can only run a handful of hunts each month because it takes so long to build a threat hunting package from scratch. With a content platform, a team can dramatically improve the mean time to deployment (MTTDp), including the number and speed of hunts, without sacrificing efficacy or lumbering itself with false positives.

# WHEN WE SAY CONTENT, WE MEAN CONTENT

Most threat hunting content lacks context and guidance. Content platforms solve this issue by providing comprehensive threat hunting packages.

## WHAT'S IN A THREAT HUNTING PACKAGE?

Packages contain everything a threat hunter needs to conduct a reliable, repeatable hunt.



### QUERY

A Pre-define A pre-configured query for security tools (e.g., SIEM, EDR) to detect adversary TTPs.



### USE CASE

A clear, concise overview of the hunting package, what it detects, and how.



### DEPLOYMENT GUIDE

Details of any package customizations or configurations needed prior to deployment



### RUNBOOK

Details of any package customizations or configurations needed prior to deployment



### REMEDIATION

Best practice remediation guidance for analysts to respond to the threat.



### EMULATION & VALIDATION

Tools to emulate threats in the hunter's environment in a non-destructive manner.



### CONTEXTUALIZATION

Full threat context, CTI, and alignment to cyber frameworks (e.g., MITRE ATT&CK).

## WHY IS REMEDIATION GUIDANCE IMPORTANT?

When they detect a threat, many organizations simply reimage the infected asset or endpoint. This is VERY dangerous.

If the infection has been present for some time, it has probably spread to other assets. Reimaging the infected asset hides the problem, eliminates forensic evidence, and warns the attacker to be more careful.

By following best practice remediation processes, incident responders can uncover the full extent of a threat and remove it from their network.

## THREAT HUNTING CONTENT PLATFORM BENEFITS

- ✓ **More and faster hunts.** It can take weeks to research, build, validate, and contextualize a hunt. A content platform provides instant access to hunting packages built by an expert team—with new packages available within hours or days of a threat surfacing.
- ✓ **No query customization required.** Content packages include queries pre-configured to the hunter's environment and tool stack. Having this significantly decreases the mean time to deploy (MTTDp) for security teams.
- ✓ **Minimizes manual analysis time.** Quality packages return a manageable number of results with few false positives.
- ✓ **Easily searchable.** Packages are enriched with relevant context and tagged with alternate names for threats and groups, making it easy for hunters to identify content that fits their needs.
- ✓ **Guides threat hunting.** The packages provided by a content platform are based on current TTPs, high-fidelity CTI, and adversary behaviors. Threat hunters can analyze the latest packages to see which threats currently focus on their industry or location and run those hunts.
- ✓ **Guides data collection.** Threat hunters aren't logging experts and don't always know what data is available. Having top-tier content can help them understand the logs their organization should collect and make a case for change if they don't have them.
- ✓ **Informs ongoing monitoring.** Once a threat hunt has run successfully, it may be suitable to operationalize as an automated detection. Not all content is suitable, but content linked to specific TTPs can often inform ongoing monitoring.

- ✔ **Aids threat hunter development.** Access to high-quality content helps threat hunters develop their skills and understanding by studying how expert hunters research and develop a hunt.
- ✔ **Helps to upskill aspiring threat hunters.** To fill the skills gap, many organizations aim to upskill top performers from other security disciplines. Access to packages developed by expert threat hunters helps aspiring hunters develop their skills.

**"IT SHOULD BE A GOAL OF ANY THREAT HUNTING ADVOCATE TO LOWER THE [ENTRY] BAR SO INEXPERIENCED ANALYSTS CAN PARTICIPATE."**

**— JOSH LIBURDI, SECURITY TECHNOLOGIST, SQRRL.**

- ✔ **MUCH lower cost than hiring more FTEs.** Hiring more threat hunters costs hundreds of thousands of dollars. A content platform can help a threat hunting team dramatically increase productivity for less than the cost of a single FTE.

## **MORE HUNTS, BETTER RESULTS, LESS TIME**

Even advanced threat hunting teams can struggle to run more than 3-4 hunts per month when producing all their own content. With a content platform, one financial institution increased its hunt output by 5X and freed up hunters to focus their efforts on the most critical hunts.

**"THE NUMBER ONE BENEFIT [OF A THREAT HUNTING CONTENT PLATFORM] IS OUR EXISTING THREAT HUNTERS ARE EXCITED TO WORK WITH THE CONTENT, AND THEY ARE INSPIRED BY IT. IT'S ALSO GREAT FROM A HIRING PERSPECTIVE BECAUSE IT GIVES NEW THREAT HUNTERS A FOUNDATION TO GET UP AND RUNNING QUICKLY."**

**— THREAT HUNTING TEAM LEAD, FORTUNE 500 FINANCIAL SERVICES INSTITUTION**



# A THREAT HUNTING CONTENT PLATFORM IN ACTION

## CYBORG'S THREAT HUNTING CONTENT PLATFORM FUELS 80% OF HUNTS FOR A LARGE US FINANCIAL INSTITUTION

### THE COMPANY

Acme Financial\* is a US-based Fortune 500 financial services institution. It has a highly mature 600-strong cybersecurity function complete with a CTI capability, established blue and red teams, and dedicated teams covering CERT, IT GRC, third-party management, and more.

\*Not their real name, *obviously*.

### THE CHALLENGE

In 2020, Acme launched a threat hunting capability. The goal was to 'deepen' proactive security capabilities and move beyond basic indicator-fueled monitoring and alerts. Starting with two FTEs, the company planned to expand to five FTEs over 24 months.

Initially, the threat hunting team faced three challenges:

1. Limited time to develop effective threat hunts—the team needed 2-3 weeks to plan and build each hunt.
2. Mapping hunts to gaps using models like MITRE ATT&CK.
3. Difficulty growing the team and bridging knowledge gaps due to lack of available education and high-quality content.

Acme needed a supplemental product to provide specialized threat hunting content and education geared to the financial services industry.

## THE SOLUTION

Acme assessed several open source and paid threat hunting content platforms. In most cases, the company was concerned about the lack of transparency provided into relevant intelligence and context for each hunt. This information was internally housed and not shared with the customer.

Another concern was that some platforms required a locally installed agent, which would further complicate Acme's tool stack.

The company agreed to a trial with Cyborg Security for specific hunting packages. In particular, Acme was impressed by Cyborg Security's:

- ▶ Continuous validation and improvement of threat hunts.
- ▶ Decay modeling to remove outdated hunts from the portal.
- ▶ Automated customization of queries to fit Acme's environment.
- ▶ Full access to relevant context, CTI, and analyst runbooks.

## THE RESULTS

Acme formalized the partnership and began using Cyborg Security's content platform for up to 80% of its hunting operations. Benefits included:

- ✔ Huge time savings—from ~3 weeks per hunt to ~1 day.
- ✔ Content transparency gives confidence that each hunt is effective.
- ✔ Helps current and prospective threat hunters build their skillset.
- ✔ Developing hunts for new threats takes days, not weeks.
- ✔ Maps to MITRE ATT&CK, Cyber Kill Chain, and Diamond Model.
- ✔ Helps Acme identify logging and data deficiencies.
- ✔ Validation and emulation tools aid in 'sense checking' content.

**"PERHAPS THE NUMBER ONE BENEFIT WE'VE SEEN FROM WORKING WITH CYBORG IS OUR THREAT HUNTERS ARE EXCITED TO WORK WITH THEIR CONTENT. IT INSPIRES THEM TO DO A GREAT JOB, AND IT'S ALSO GREAT FROM A HIRING PERSPECTIVE BECAUSE IT HELPS NEW HUNTERS GET UP AND RUNNING FAST IN OUR INDUSTRY."**

**— THREAT HUNTING TEAM LEAD, ACME FINANCIAL\***

*\*Again, not their real name.*

# A BUYER'S GUIDE TO CONTENT PLATFORMS

Would your threat hunting team benefit from a powerful content platform? Be careful—they don't all provide the same quality of content. Before you buy, make sure your preferred solution fits these criteria.

---

## EXHAUSTIVE CONTENT DEVELOPED BY 'BEST OF THE BEST' HUNTERS.

Why accept anything less? In addition to a pre-configured and fully optimized query, every threat hunting package should include:

- ✓ Human-readable documentation for the threat hunter.
- ✓ A deployment guide for any necessary customizations.
- ✓ Tools to emulate threats in the hunter's environment in a non-destructive manner.
- ✓ Full threat context, including related groups, motivations, behaviors, and alignment to cyber frameworks.
- ✓ Runbooks and remediation guidance for security analysts.

---

## INFORMED BY CURRENT, VERIFIABLE, AND TRANSPARENT THREAT RESEARCH.

Intelligence feeds operations, and poor CTI leads to ineffective threat hunts. The platform provider should create content based on exceptional research capabilities—*and* provide you with total visibility, making it easy for hunters to 'pivot' from a threat hunting package to relevant research.

---

## RIGOROUSLY VETTED, TESTED, AND EASY TO VALIDATE.

Unvetted content wastes precious time and resources. The platform should only provide content that is current and high-fidelity. Each package should include verification and emulation tools that enable threat hunters to easily test each piece of content for efficacy.

---

## TOOL AGNOSTIC AND OPTIMIZED FOR YOUR ENVIRONMENT.

Altering queries manually to fit your environment and tool stack is time consuming. The platform should automatically scope your environment using metadata and instantly customize each query to your needs.

### HOW DO I EVALUATE QUERY OPTIMIZATION?

A well-tailored query runs quickly and returns few false positives. Outputs are mapped to your environment, with accurate naming conventions for fields and indices.

---

## FULLY TAGGED AND EASILY SEARCHABLE.

Packages should be right at the hunter's fingertips. All packages should be fully tagged with details of the threat, alternate names, related groups, impact and priority, in-the-wild sightings, and affected asset groups.

---

## PROVIDES **ONLY** "LIVING" CONTENT.

Most sources of threat hunting content are overwhelming, and most entries stale and useless. The platform should use decay modeling to identify and update non-current content.

Decay modeling is not limited to threat hunting packages. The platform should also 'decay' outdated intelligence on related actors, behaviors, etc.

---

## **MAPPED TO MITRE ATT&CK AND OTHER CYBERSECURITY MODELS.**

Every hypothesis should focus on a high-risk TTP that isn't covered by other security capabilities. This is the core of risk-based hunting.

To support this, the platform should map content to popular models like ATT&CK, Cyber Kill Chain, and Diamond Model for Intrusion Analysis.

---

## **NO LOCAL CLIENT NECESSARY.**

Security teams have enough tools and complexity to handle. There's no reason why a threat hunting platform should require a local client—a web-based portal gives threat hunters instant access to content without adding complexity to the technology stack.

# THE INEVITABLE PITCH

In this paper, we wanted to highlight the value of structured threat hunting, the challenges hunting teams face, and—in our humble opinion—the best way to solve those challenges. We’ve tried to do that in a fair and unbiased manner without calling out vendors and cybersecurity professionals that are doing their best to make the digital world a safer place for all of us.

We also tried to avoid unnecessary fear-based marketing nonsense and pictures of hoodie-wearing hackers.

(We had to convince our marketing department about that last part. You’re welcome.)

At the end of the day, though, we do have something to sell: a threat hunting content platform. And we believe it’s the best darn threat hunting content platform out there.

Our team of dedicated threat hunters is second to none. They develop complete threat hunting packages for brand new threats in just a few hours or days. We also have a dedicated CTI team that focuses exclusively on intelligence that fuels effective threat hunts—and you get total visibility of that intelligence in every package.

So that’s it. If you like our approach, and you’d like to boost your threat hunting capabilities by as much as 5X—for **less than the cost of a single FTE**—you can arrange a free, no obligation trial of our HUNTER platform at the link below.

## SEE HUNTER IN ACTION

Get started hunting for FREE! Get a  
Community Account on the HUNTER  
Platform at [cyborgsecurity.com](https://cyborgsecurity.com)!



[in](#) Cyborg-Security [t](#) @CyborgSecInc [v](#) Cyborg Security  
[f](#) Cyborg Security [i](#) @CyborgSecInc [d](#) Cyborg Security