

The Preemptive Cybersecurity *Handbook*



• FROM SEARCHLIGHT CYBER •

A practical guide to find and fix
what attackers will look for next

The Preemptive Cybersecurity *Handbook*

A practical guide to find and fix
what attackers will look for next



Searchlight
Cyber

| Contents

Foreword	6
----------	---

PART ONE: WHY? 10

What Does it Mean to be Truly Preemptive in Cybersecurity?... And Why is Everyone Suddenly Talking About it Now?	12
---	----

A Landscape That Refuses to Stand Still	18
---	----

Zero Days and the Attacker's Perspective	24
--	----

The Future of Zero-Days: The cPanel Case Study	32
--	----

The Six Principles of Preemptive Security	39
---	----

PART TWO: WHAT? 46

Introducing Preemptive Threat Exposure Management	48
---	----

PTEM Foundations: Know Your True Attack Surface	54
---	----

Attacker Reality: The Threat Intelligence Layer	66
---	----

Ransomware: Where Preemptive Defense Matters Most	81
---	----

PTEM Visualized	92
-----------------	----

PART THREE: HOW?	94
Building a Preemptive Organization	96
Why Today's Cybersecurity Success Metrics Miss the Point	109
The Preemptive Cybersecurity Maturity Model	118
Leading the Transformation	130
PART FOUR: REFERENCE & LEARNING	134
The Practitioner's Guide to the Dark Web	136
Hacking Forums	140
Ransomware Leak Sites	153
Dark Web Marketplaces	164
PTEM Buyers Guide	176
The Six Capabilities Every PTEM Platform Should Provide	180
Glossary of Key Terms	196

Foreword

Michael Gianarakis, CEO, Searchlight Cyber

Some of the most important ideas in cybersecurity don't begin in the boardroom, but in the messy reality of trying to find something before everyone else does. That was the world I came from: a world shaped by offensive security and bug bounty work. I quickly understood that strong security posture was much more than a checklist of tools, it was a discipline built around speed, scale, and understanding what an attacker can actually see.

When I and my co-founding colleague Shubham Shah (Shubs) founded Assetnote, the Attack Surface Management company which is now part of Searchlight Cyber, the main challenge facing organizations was a lack of visibility over their true attack surface. Though things have moved on, the overarching issue is one that still persists: the problem was moving faster than the industry's ability to keep up. Attack surfaces were expanding continuously, and organizations' estimation of their attack surface versus the true extent of what an adversary could actually reach and exploit, were worlds apart. We built a tool firstly to help us find vulnerabilities faster, but gradually saw how we could close the gap for organizations. And, in doing so, we learned something that has stayed with me ever since: visibility by itself is not enough unless it leads to action.

That insight is one of the reasons I believe so strongly in preemptive cybersecurity as a way of thinking about the future of our industry, and Preemptive Threat Exposure Management as the practice to achieve that future. Security has spent too long being defined by response. Important as response is, it still accepts an attacker getting there first. Preemptive cybersecurity aims to answer a much better question: how do we identify exposure, context, and threat activity early enough to change the outcome altogether? That is the shift the ideas and frameworks in this book aim to accelerate.


Over the years, I've seen the market slowly catch up to ideas that once felt obvious only to those living closest to the problem. I've also seen how easy it is for the meaning of Attack Surface Management to become diluted when it is reduced to a category label instead of a practical capability. The real work is harder than that. It requires continuous discovery, a focus on real threats rather than noise, and the ability to connect external exposure with the threat signals that show where risk is becoming real. The focus should be on the ultimate outcome of reducing exploitable exposure before attackers act, and achieving this requires capabilities that enable confidence within the team executing the work.

This book is for the security leaders and practitioners who know that the old model is not enough. It will help you think more clearly about your exposure, understand how PTEM fits into a modern security program, and build a more disciplined, more proactive approach to cyber risk. Most importantly, it will help you move from asking what happened after the fact to understanding what is coming next, and what you can do before it becomes an incident.

I'm proud of the journey that brought us here; from the early days of Assetnote with Shubs, to the broader vision we now share at Searchlight Cyber. But more than that, I'm excited for what this book can do for the people reading it: sharpen their thinking, strengthen their programs, and help them get ahead of threats rather than merely reacting to them.

Foreword





What Does it Mean to be Truly Preemptive in Cybersecurity?... And Why is Everyone Suddenly Talking About it Now?

Pre-emptive

If something is pre-emptive, it is done before other people can act, especially to prevent them from doing something else.

Cambridge English Dictionary

The idea that you can deter or prevent an anticipated situation isn't some modern cyber-revelation. It's been a bedrock of human strategy for millennia, from ancient battlefield maneuvers to the basic tenets of every successful logistics chain in history. With the right insights and the right tactics, the power to head off trouble was a mark of being well-informed and well-prepared. So, if we've understood the value of the preemptive strike since the dawn of organized conflict, why does it feel like such a moonshot for modern cybersecurity? Why is it that in every other high-stakes field, being truly 'preemptive' is a sign of a professional, yet still appears to be so out of reach for many organizations' security strategies?

The answer lies in the fact that we've been trying to win a game of foresight using the tools of hindsight.

For a long time, the industry standard for "security" has been firmly rooted in the reactive. We've been trained to build a perimeter, watch the gates, and wait for the alarm to sound. When it does, we jump into action, contain the threat, and call it a win. For years, that "detect and respond" rhythm was enough. Back when infrastructure moved at the speed of a human change request, we could afford to be the ones waiting. We had the luxury of a stable, predictable environment, and our adversaries, for the most part, were playing by the same slow-motion rules.

But the ground has shifted under us. Today, the perimeter is no longer as tangible, and it certainly couldn't be called static; it's a sprawling, fragmented, and constantly evolving digital ecosystem. While we've been busy perfecting our ability to respond, the threats have become automated, persistent, and AI-enabled. Responding was never going to contain the damage. So what do we do? Treat an eventual breach as inevitable and rely on our ability to weather the storm? Let's not be defeatist.

The good news is that being able to preemptively stop threats is not just an aspiration. After all, you have this book in your hands, likely looking for a tangible way forward. By leaning into better intelligence, processes, and making your environments more dynamic, you can actually take the initiative. Being preemptive is about using insight to stay three steps ahead of the trouble, turning security from a frantic game of catch-up into a strategic foundation for everything else the business wants to build.

The Illusion of Control: Why We Ended Up Reactive

Cybersecurity didn't start out trying to be reactive. But traditionally, organizations had a perimeter, a patch cycle, and for the most part that meant a reactive model of "scan, detect, alert, patch" was a good fit. You could afford to be the last one to know because the environment changed about as often as the seasons.

But then, the business world decided it wanted to move at the speed of the cloud.

That single, sturdy castle needed to be traded in for a sprawling, global grid of SaaS apps, microservices, and third-party dependencies that update every few minutes. Rapidly, the perimeter dissolved. While you were busy trying to map this new reality using the same old point-in-time audit mindset, your adversaries were busy doing what they do best, with far fewer restrictions, finding the gaps in your infrastructure that you didn't even know existed.

The result is a whack-a-mole cycle. Organizations don't end up reactive because of lack of talent, but because they've been trying to solve a dynamic, high-speed problem with a toolkit designed for the 2010s. There's an obsession with dwell time: the amount of time an attacker sits in our network before they're kicked out. But dwelling on dwell time is a symptom of a deeper issue: we've organized our entire industry around the incident. We treat the breach as the event, and the response as the solution.

Cybersecurity should have always been preemptive. It's the difference between guessing where you're vulnerable, and actually knowing what's going to happen in advance. By defining security success as responding to what's already happened, you're just ensuring you'll always be the second person to act in any given scenario.

Attackers are going to target your organization. They probably are already. That's just an unfortunate fact of the world we live in. But it's possible to break out of the reactive cycle. This requires viewing and operationalizing security as what it should have been all along: an active, continuous process of disappointing attackers at every step.


“
*we've been
trying to win
a game of
foresight using
the tools of
hindsight.*

The Hottest Ticket in Town

We're not going to pretend that we've come up with the concept of preemptive cybersecurity; everyone's talking about it. But if it's such a timeless principle, why now? The truth is, we've reached a point where the status quo is simply unsustainable.

We're seeing this pivot because, for the first time, the technology is actually catching up to the ambition. We have the AI to model attacks, the simulation engines to test our own defenses, and the intelligence platforms to spot the bad actors long before they get within striking distance. Creating an environment where attacks fail to gain traction is no longer a hypothetical dream.

The industry is realizing that resilience isn't found in the wake of an incident, but in the deliberate, preemptive steps before the incident happened.



A Landscape That Refuses to Stand Still

For years, the vulnerability disclosure cycle offered organizations a degree of breathing room. A flaw would be identified, a CVE would be published, and somewhere between days and weeks later, security teams would have patched the affected systems and moved on. That window, imperfect as it was, made reactive security viable. But that window is now measured in hours.

At the time of writing this book, The Zero Day Clock (zerodayclock.com), a resource developed by Sergej Epp of Sysdig, is tracking Mean Time to Exploit (TTE) at 10hrs in 2026. TTE measures the gap between CVE disclosure and confirmed exploitation, and to put this into perspective, Mean TTE for weaponized exploits was 2.3 years in 2018. Even just two years ago in 2024, organizations had on average 56 days between disclosure and exploitation.

Also worth highlighting is the prevalence of zero-day vulnerabilities; the percentage of exploited CVEs where exploitation occurred before or on the day of disclosure. In 2025, we crossed the threshold where the majority of CVEs exploited in the wild were zero days, when that figure stood at only 16% in 2018. At the time of writing, 74.4% of exploited CVEs in 2026 were classified as zero days. For many years, the argument in the industry has been that the vast majority of breaches are caused by other attack vectors, so security professionals should

not be blinkered into focusing on so-called 'sensationalized, unpreventable threats' but rather focus on core security hygiene. That's fair advice, but the numbers shown here highlight how quickly both known and unknown vulnerabilities are successfully weaponized by threat actors.

Threat actors have built entire operational infrastructures designed to weaponize vulnerabilities at speed. The moment a proof-of-concept is published, exploit code circulates in criminal forums. Initial access brokers scan the internet for exposed systems before most organizations have even assessed whether they are affected. In some of the most consequential attacks of recent years, exploitation began before a public advisory had even been issued. Threat actors discovered or purchased vulnerabilities that defenders did not yet know they needed to worry about.

The numbers show more than a marginal change in attacker speed. What we're seeing is the elimination of the reactive window altogether. Patching cycles, change management processes, and vulnerability triage workflows were designed for a world where time existed to use them. In too many cases, that time no longer exists.

“ *What we're seeing is the elimination of the reactive window altogether.* ”

Adversaries with Very Few Limitations

If shrinking exploitation timelines were the only challenge, defenders would at least know what they were working against. But the nature of the adversary itself has changed.

The industrialization of cybercrime over the past decade has turned what were once isolated threat actors into specialized ecosystems. Ransomware groups purchase initial access from brokers. Exploit developers sell their work to multiple customers. Affiliates operate off-the-shelf toolkits with professional support structures. The entry barrier has dropped dramatically, while the ceiling of what even mid-tier attackers can achieve has risen sharply.

Now, artificial intelligence is tilting the balance further. Attackers are using AI to scan exposed infrastructure at a scale and speed that manual reconnaissance could never match, identifying misconfigured services, vulnerable dependencies, and exploitable paths across thousands of targets simultaneously. They are using it to generate convincing phishing content personalized at scale. And increasingly they are using it to identify novel vulnerabilities and write the exploit code needed to breach even the most well-defended organizations.

None of this means that AI has made attackers invincible. But it has made them faster, more scalable, and capable of operating at a level of sophistication that would previously have required significant expertise and resources. The same capabilities that make AI a useful tool for defenders are resulting in an inevitable arms race, and the rate of discovery of new vulnerabilities is not going to slow down any time soon.

The Third Party Trust Chain

At the same time, attackers have become more strategic about where they direct that capability. A hardened enterprise perimeter is a difficult and expensive target. A third-party supplier with access to that enterprise, operating with fewer resources and less scrutiny, is a much easier one. Supply chain and third-party compromises have become one of the defining attack vectors of this era, because the boundaries of what constitutes your attack surface have expanded far beyond organizations' direct control.

The trust chain is what makes third-party software vulnerabilities so dangerous: when an organization relies on external vendors, open-source libraries, SaaS platforms, or managed services, it is effectively extending trust beyond its own walls, and attackers only need to compromise one link in that chain to create downstream impact across many customers at once. A flaw in a widely used component is rarely just a technical issue in isolation; it becomes a business risk because it can expose data, disrupt operations, or provide a foothold into environments that the organization does not directly control. That is why security teams need to think less about individual vulnerabilities in isolation and more about the trust relationships that turn those vulnerabilities into genuine attack paths.

An Attack Surface Area That Constantly Expands

One of the most enduring challenges outside of attacker speed and sophistication is the sheer scale of what defenders are now responsible for protecting.

As we've outlined, the modern attack surface is not a perimeter. It is not a defined set of assets inside a network boundary. It is a constantly shifting collection of cloud infrastructure, SaaS applications, APIs, third-party integrations, remote endpoints, shadow IT, acquired companies and their legacy environments, and internet-facing assets that in many cases organizations are not fully aware they possess.

The lack of visibility itself is bad enough, but what it creates secondarily is a prioritization issue.

When every scan returns thousands of findings, every tool produces its own feed of alerts and vulnerabilities, and the list of things that could theoretically be addressed never gets shorter, the signal that actually matters begins to disappear into the noise. A genuinely critical exposure, an internet-facing asset actively being targeted by threat actors, can sit buried beneath hundreds of lower-risk findings simply because no system exists to understand its real-world significance. Effort gets spent, but risk does not get reduced.

This is the state in which many security teams now operate. Not through a lack of capability, but because the model they are working within was not designed for an environment of this scale and complexity.

What This Means for Detect and Respond

Each of these forces; collapsing exploitation timelines, increasingly capable and scalable adversaries, and an attack surface that is growing faster than it can be governed, would be challenging in isolation. Together, they create a compounding problem.

Detect and respond assumes that compromise is an event with a beginning you can observe and a timeline you can act within. But when exploitation can begin before you are aware of a vulnerability, when attackers can move laterally and establish persistence quickly, and when the sheer volume of alerts makes it harder to identify what is real and urgent, breach detection loses all of its value as a first line of defence.

Detection and response absolutely have a role to play. When a breach occurs, the ability to identify and contain it quickly is critical. But as a primary security strategy, as the answer to the question of how an organization keeps itself safe, it is no longer sufficient on its own.

The forces reshaping the threat landscape are not temporary conditions but structural shifts. And they demand a structural response.

Zero Days and the Attacker's Perspective

Shubham Shah

Chief Security Research Officer, Searchlight Cyber

The Zero-Day Clock: Why Time Is Against Us

Zero days are about more than exotic bugs. Really, it all comes down to timing. In practice, a zero-day is any vulnerability an attacker can use before defenders have patches, signatures, or even a shared name for it. The defining feature is that defenders don't have a playbook for it yet.

Over the last few years, we've watched the "zero-day clock" move sharply against defenders. In our work and across industry data, we see a growing share of exploited vulnerabilities being hit on or before the day they're disclosed. In other words, by the time a CVE is assigned and an advisory is written, exploitation is already under way or not far from it. The window between someone being aware of a bug and the bug being used against real organizations has compressed down to days or even hours.

Seeing First: Origins in Bug Bounty

Before I founded Assetnote, I faced a problem. I was losing bug bounties because other researchers were seeing assets before I did. They weren't necessarily smarter or more creative; they just had better reconnaissance. They knew about new subdomains and new internet-facing services before I did, and in this game you only get paid for first place.

So I started building my own tooling. The first version of Assetnote was something I made myself and presented at B-Sides Canberra. The idea was basically that if we could continuously discover and monitor subdomains and other external assets as they appeared, we'd get notified about exposed infrastructure before it showed up in anyone's weekly scan or spreadsheet.

That's where the idea of "ephemeral vulnerabilities" really clicked for me. A lot of the most interesting bugs live on assets that only exist for a short time. If your visibility runs on a weekly cadence, you'll never see those. If you have a system that's watching continuously and telling you about new assets within minutes or hours, you catch them while they're still alive.

One of the early wins that cemented this approach was against Slack. Our tooling picked up staging subdomains that briefly exposed Git repositories. Pulling them down gave us access to core application source code and hard-coded secrets.

That led to what their CISO later described as one of the largest external bugs they'd seen.

In the bug bounty world, that kind of advantage is key to getting paid. In the enterprise world, it's the opportunity to quietly fix a critical issue on a short-lived asset and before it's found by a threat actor. Assetnote grew out of the realization that the same advantage that helps a single hacker win a bug bounty can be turned into an organizational capability.

“

The real core output of offensive security research should be time advantage for defenders.

From Hacker Edge to Preemptive Framework

We took the patterns that worked in bug bounty recon: continuous asset discovery, change detection, early alerts on new internet-facing services, and wired them into a platform. Many tools we made available open-source for the wider benefit of the community. We made sure that every interesting finding became a potential rule, not just a story. If we noticed a specific SaaS vendor's customer subdomains following a pattern, that pattern went into our discovery engine so it would fire automatically for every customer.

That's essentially what preemptive cybersecurity is: the organizationalization of an offensive mindset. It's taking the way good hackers see an attack surface and getting it running 24/7. Recon feeds research. Research feeds visibility. Visibility feeds prioritization. And the whole loop is designed to give defenders the same "see first, act first" advantage we built for ourselves in bug bounty.

What Scanners and Advisories Cannot See

Most organizations already run scanners and consume advisories. They patch, track CVEs, and follow best practices. None of those are inherently wrong; but they have a built-in blind spot.

Many of the most impactful bugs we've found do not fit neatly into what a scanner is looking for. They could be the result of a specific way two services are linked together, or a combination of default settings that, when combined, open a path nobody anticipated.

No generic scanner was written to look for that because it's not "a vulnerability" in isolation; it's a chain that only exists in that context. Scanners with basic CVE-matching only focus on the 'known' rather than the as yet unknown, and without context also produce significant false-positives.

Advisories can only talk about what has already been found, named, and documented. Until a vulnerability crosses that threshold, it doesn't exist in their world. Research exists to fill that gap. Our job as offensive security researchers is to go looking in the places the tools can't see yet and bring back enough detail that vendors, defenders, and eventually the wider ecosystem can act.

Research as Time Advantage, Not Just Bug Hunting

From the outside, research can look like a game of "who found the coolest bug." There are bug bounty leaderboards, conference talks, social media posts, it's easy to think the main output is bragging rights.

But the real core output of offensive security research should be time advantage for defenders.

Every novel vulnerability that is discovered, validated, and fixed before exploitation buys defenders a window. That might be a few days. It might be a few weeks. Sometimes it's longer. In that window, the organizations connected to the research can either patch or apply rapid mitigations before the patch exists and the wider industry scramble begins.

We've seen this play out in practice. When we responsibly reported certain issues to large vendors, they quietly developed patches and guidance. In parallel, we built detections and mitigations that we could deploy through our platform. By the time the advisory went public, our customers had already moved. They weren't racing the internet; they were cleaning up the last edges of a problem they'd been mitigating for days.

Humans in an AI-Heavy World

Over the last couple of years, we've been experimenting with frontier AI models in our research workflow. We've used them to help review code and to explore parts of large, messy attack surfaces more quickly than a human could on their own.

Used well, these models absolutely make us faster. They help us stage work, explore more candidates in parallel, and get to "the interesting bits" sooner. Some of our recent disclosures to large enterprise vendors were accelerated because a model helped us sift through boilerplate code and focus our attention on the parts of the codebase or configuration that actually mattered.

Used badly, however, they create a lot of noise. We've seen disclosure programs flooded with AI-generated reports describing issues that aren't exploitable in any realistic way, or combining unrelated behaviors into imaginary "chains." Even experienced operators can fall into the trap of chasing down something that only exists in the model's head.

The best way to look at AI is as a force multiplier for people who already know what they're doing, not a substitute for judgment. A model can suggest an interesting path. It can substantially speed up the real grind work. It can help you explore "what if" scenarios. And overall it changes the role of the researcher in terms of where they point that capability in order to find the bugs that really matter.

For all the hype around AI, vulnerability research is one area that AI has pushed into a genuinely new era. My team and I weave it into existing workflows where it makes sense: triaging large inputs, testing variations, generating candidate payloads. But the human job doesn't go away. It shifts more toward context: deciding where to point the models, filtering their output, and turning their suggestions into real, grounded security improvements.

Relevance Over Novelty

There's always a temptation to chase novelty for its own sake; obscure software, extremely constrained exploit scenarios, bugs that make for great talks but affect five deployments worldwide. Sometimes they're necessary for pushing the field forward, but they don't always move the needle for the organizations we're trying to protect.

Effective security research has a bias toward relevance. We ask: what are the technologies our customers actually run? Which platforms show up over and over again in their attack surfaces? Which services are already popular targets for attackers, or sit in the critical path of business operations?

That's how we ended up spending serious time on things like cPanel, which we'll come to in more detail in the next chapter. We looked into it because we see it everywhere: in hosting environments, exposed to the internet. If an attacker compromises something like that, they could inherit control over a huge slice of your environment.

From a preemptive standpoint, there is an obvious choice. If we can find one serious vulnerability in software that's everywhere, and responsibly get it fixed and mitigated, we change the risk profile for a lot of people at once.

Zero-Day Research at the Heart of Preemptive Security

Preemptive security is all about winning back time.

We know exploitation timelines are shrinking. We know attackers are increasingly hitting vulnerabilities before or on the day they're disclosed. We know that scanners and advisories, while important, only see what has already been named and catalogued.

Zero-day research is how we push back. By combining deep recon, focused research, and tight feedback loops into our tooling and processes, we can create windows where defenders move before attackers, not after. We can surface the exposures that would never show up in a scanner report, and aim our effort at the parts of the stack that actually matter in real environments.

The Future of Zero-Days: The cPanel Case Study

The vulnerabilities we discovered in cPanel (in early 2026), a platform sitting quietly behind millions of websites and mailboxes worldwide, serve as a perfect case study of how AI has changed research: even when you run a proactive research program and lean on the latest models, threat actors can and do get there first. That collision in cPanel tells us a lot about the new era of vulnerability research and what being preemptive in cybersecurity really means.

The cPanel Wake-up Call

If you run anything on the internet in 2026, there's a good chance cPanel is somewhere on your attack surface. The technology is just so pervasive that in many cases you might not even know that some of your domains have it exposed.

Over the course of the first few months of 2026, my research team got to work pulling apart the cPanel stack and uncovering issues that cut right to the heart of internet infrastructure. The much-publicized CVE-2026-41940, the authentication bypass vulnerability in cPanel & WHM, allowed an unauthenticated remote attacker to gain administrative access to the control panel and, by extension, the underlying host. CVE-2026-29205 allowed access to sensitive files with root-level privileges, which could give an attacker extensive access to files across the vast cPanel footprint.

We discovered these issues in early April, but in the case of CVE-2026-41940, not ahead of the attackers. Public reporting suggests that threat actor exploitation of the authentication bypass may have started as early as late February, roughly ten weeks before many organizations even heard of the CVE identifier.

In other words, even with a proactive research program and access to frontier AI models to accelerate the work, threat actors had already been living in this zero-day for weeks. That delta between first exploitation and industry awareness is the gap that now defines modern vulnerability management.

A Knife Fight in a Phone Booth

Vulnerability research has always been brutally competitive. Historically, research teams performed their work through manual investigation and close interrogation of source code. Now they can leverage these models in a highly targeted way, allowing them to perform deeper analysis across much larger datasets and codebases. Tasks that previously took days of focused human review can now be staged and parallelized across model-assisted analysis. The limiting factor is now intent rather than tooling.

In our research into cPanel, we benefited from this acceleration. Frontier models helped us move faster through reconnaissance and code analysis, and contributed to uncovering the multiple vulnerabilities. But the truth is that the adversaries have similar tools and fewer constraints. They only need to be first once.

The bar to compete in vulnerability discovery has dropped, but to be first to find and to meaningfully reduce risk, the stakes have never been higher.

Reactive Patching is Dead: The Exposure Window Starts Long Before CISA KEV

For years, security programs have optimized around reactive vulnerability management. Waiting for a CVE, watching it climb on the CISA Known Exploited Vulnerabilities (KEV) list, scramble to patch.

That model assumes a world where there is a comfortable buffer between “vulnerability exists”, “exploit exists”, and “the defender knows and can act.” In the age of AI-accelerated research, that buffer is gone.

These bugs didn’t suddenly become dangerous the moment an advisory was released or researchers reverse-engineered the patch. It was dangerous the entire time it existed in production software, including the period when nobody on the defender side had written a single detection for it.

Zero-days are no longer rare events, but a constant background condition. The only thing that changes is who finds them first and how long they are quietly exploited before the rest of the industry catches up.

If your definition of “preemptive” is moving quickly once something hits CISA KEV, you’re fighting a losing battle.

What Preemptive Should Mean Now

Much of the industry language around “preemptive” or “proactive” security still centers on reverse-engineering n-day vulnerabilities once they’re public: pulling apart proof-of-concepts, writing detections for existing exploits, and retrofitting controls around advisories. That’s not enough.

In this new era, preemptive has to mean discovering and mitigating vulnerabilities before threat actors do, particularly in the software that underpins the internet’s critical infrastructure.

Preemptive cybersecurity must be underpinned by systematic, hypothesis-driven zero-day research into critical enterprise platforms, and detection strategies that don’t rely on CVE identifiers, but exposure context, exploitability, and attacker behavior signals.

Everyone is Equally Elevated at Finding Vulnerabilities


Frontier models are continually improving, and threat actors will not slow down in their efforts to find and exploit novel vulnerabilities. Therefore, security strategies need to be designed around a few harsh realities:

You cannot assume the defender will always find a bug first. In many cases, the opposite will be true. The goal is simply to reduce your exposure window as much as possible. You cannot rely on public disclosure as your starting gun. By the time a vulnerability has a CVE, an advisory, or a KEV entry, you should assume that exploitation is already well underway.

Novel vulnerability research is a core defensive function, not a vanity project. People whose job it is to break your own (and the internet's) critical dependencies, augmented by AI, not replaced by it, are more important than ever.

You need visibility and telemetry that makes unknown-unknowns observable. If you cannot see what you're exposed to, you will only ever learn about zero-days from someone else's incident report.

What differentiates defenders is what you do with the head start you still have in visibility, context, and intent.



The Real Time Era of Cybersecurity

Finding vulnerabilities in cPanel has been super interesting. Despite having a significant head start compared to our industry peers in analyzing cPanel for security issues, there were clearly threat actors that had a similar idea to us when it came to auditing internet critical software proactively. For us, this is the beginning of the new age. Everyone is equally as elevated towards finding and fixing vulnerabilities.

The vulnerabilities in cPanel didn't suddenly become important when they hit an advisory, or when they appeared in a feed, or when someone wrote a catchy blog title about "the internet falling down." They were important the entire time they existed, long before any of us gave them a name.

Organizations' job now is to build security programs that behave as if those bugs are already out there and already being hunted with the best tools available. Because they are.

The Six Principles of Preemptive Security



Assume You're Exposed. Focus on Reducing It.

Preemptive cybersecurity requires acknowledging you are already exposed in ways you don't fully understand.

Exposure reduction is continuous. The goal is to reduce exploitable conditions across your attack surface before attackers act.

This reframes security from: "Are we secure?" to "Where are we exposed, and how quickly can we reduce it?"



Prioritize What Matters to the Attacker, Not What's Easiest to Fix.

Not all exposures are equal, but most security programs still treat them that way.

Preemptive cybersecurity aligns prioritization with real-world attacker behavior:

- Is it exploitable in practice?
- Is it reachable from the outside?
- Does it lead somewhere valuable?

This is the shift from checkbox security to adversary-informed decision-making.



Continuously Discover. Because Your Attack Surface Won't Sit Still.

You can't fix what you don't know exists, and in modern environments, what exists is constantly changing.

Preemptive security assumes:

- Assets appear and disappear constantly
- Exposure is dynamic, not periodic

This principle pushes you away from point-in-time assessments toward continuous discovery and remediation.



Validate Security Through the Lens of Exploitability.

Controls, policies, and configurations don't matter if they don't hold up under real-world conditions.

Preemptive cybersecurity emphasizes validation:

- Can this actually be exploited?
- What would it take for an attacker to chain this into impact?
- Are our defenses observable and bypassable?

This is the difference between being secure in theory and secure in practice.



Measure Success by Risk Reduction, Not Activity.

Traditional programs measure effort: alerts handled, patches applied, tickets closed.

Preemptive cybersecurity measures outcomes:

- Fewer exposed attack paths
- Fewer exploitable entry points
- Faster reduction of high-risk exposures

This is what makes security measurable in terms the business actually cares about: risk, not activity.



Reject Inevitability. Preemption Is Possible.

“Assume breach” was a useful reset, but don’t let it drive a passive approach.


Preemptive cybersecurity rejects the idea that compromise is simply the cost of doing business. While resilience and recovery are essential, they are not the goal. The goal is to stop attacks before they succeed.



The reality is

- Most attacks follow observable paths
- Many exposures are preventable
- The capability now exists to reduce them in advance

Not every attack can be stopped, but organizations are capable of preventing far more than they realize.

Understanding these principles is the first step. Applying them is where the real transformation begins. To make this real, you need a way to continuously find, prioritize, and reduce the exposures that actually matter. In the next section, we'll break down how preemptive principles can be operationalized in a Preemptive Threat Exposure Management (PTEM) approach.





Introducing Preemptive Threat Exposure Management

The six principles in the previous chapter describe how to think about security in a preemptive posture. But principles without a practical framework are just philosophy, and you're probably thinking: "it all makes sense, but what do I do now?" In this section, we'll introduce the operational model that brings the preemptive principles to life: Preemptive Threat Exposure Management, or PTEM.

The Evolution of Exposure Management

To understand PTEM, it helps to trace the short lineage of how the industry got here. Warning: Acronyms Ahead. But we promise, it'll all make sense soon.

Gartner's research on the Continuous Threat Exposure Management (CTEM) framework in 2022¹ was a key moment that generated a lot of buzz. Built around five stages: Scoping, Discovery, Prioritization, Validation, and Mobilization—CTEM introduced many organizations to the idea that exposure management should be continuous and validation-driven, focused on what is actually exploitable rather than what is theoretically risky. Gartner's research estimated organizations aligning to CTEM were three times less likely to suffer a breach.

PEM (Preemptive Exposure Management), emerging again from Gartner in June 2025, evolved this further², introducing prediction and prevention as core objectives. Rather than continuously validating exposure, PEM adds the capability to simulate attacker behavior; modelling attack paths, running automated adversarial testing, and identifying which routes to critical assets are most realistic, before an attacker gets the chance to walk them.

PTEM goes one step further still by answering the most operationally urgent question of all: not just could an attacker exploit this, but is an attacker actually targeting this right now?


¹ Gartner, "Implement a Continuous Threat Exposure Management (CTEM) Program," Jeremy D'Hoinne, Pete Shoard and Mitchell Schneider, 21 July 2022.

² Gartner, "Emerging Tech: The Future of Exposure Management Is Preemptive," Elizabeth Kim and Apoorva Chhabra, 25 June 2025.

What Makes PTEM Different

Most exposure management approaches build a model of attacker capability. They ask: given what we know about our environment, what could be exploited?

PTEM incorporates attacker intent; the actual, observable activity of threat actors in the environments where they operate. Cybercriminals conduct reconnaissance, discuss targets, sell stolen credentials, and signal intent before campaigns are launched. When that activity is visible and feeds directly into exposure prioritization, the nature of the decisions security teams make changes fundamentally.



The difference in practice is significant. A vulnerability flagged as high-severity based on exploitability is one thing. That same vulnerability elevated because a threat actor has shown a specific targeting intent against your organization is something else entirely. One is a finding on a list. The other is the most important thing your team works on today.

This is the operational reality at the heart of a PTEM approach: exposure management that is not just continuous and validated, but genuinely intelligence-driven. Grounded in what real adversaries are actually doing, not what they theoretically could do.

The Three Pillars of PTEM

At its core, PTEM works through three integrated capabilities:

1 Continuous Attack Surface Visibility

Know what your external exposure looks like at any given moment, including shadow IT, third-party dependencies, and assets that may have been forgotten but remain active and reachable.

2 Real-World Threat Intelligence

Monitor the channels where attacker activity is observable: criminal marketplaces, closed forums, dark web infrastructure, and the signals that indicate an organization is being actively targeted.

3 Intelligence-Driven Prioritization and Preemptive Action

Bring visibility and intelligence together so that remediation effort is focused not on the longest list of vulnerabilities, but on the smallest set of exposures that represent the most credible, imminent risk.

The Standard Is Shifting

Sure, you're reading a book by Searchlight Cyber, but the emergence of PTEM reflects a genuine industry consensus that has been building across analysts, practitioners, and security leaders. It's part of a wider direction of travel, away from periodic, asset-centric vulnerability management and toward continuous, adversary-informed, intelligence-driven exposure management.

The chapters that follow break down each component of PTEM in detail; how attack surface management works in practice, what threat intelligence looks like as an operational input, and how you can build a preemptive security program that is measurable, scalable, and grounded in the real threat environment.

PTEM Foundations: Know Your True Attack Surface

Picture this: a security team arrives at a recently breached organization. The company is large, well-resourced, and, on paper, mature. They have a vulnerability management team. They have dashboards. And yet, sitting across from the incident responders, the question that no one can confidently answer is a big one: what are we actually trying to protect?

This is not an unusual situation; it's all too common. Most organizations are defending an imaginary version of their attack surface: the one they believe exists, based on the assets they know about, cataloged at a point in time that has long since passed. The real attack surface is alive and constantly expanding. It is almost always significantly larger than anyone expects. And it contains the assets that attackers will find first.

Understanding your true attack surface is the foundational pillar of Preemptive Threat Exposure Management. Knowing what you own, where it lives, what it exposes, and how it looks from the outside. Everything else in the PTEM model depends on it. You cannot prioritize defending what you have not yet discovered. And that age-old saying again: you cannot defend what you do not know exists.

The Attack Surface Problem

To understand why true Attack Surface Management matters, it helps to start with that almost embarrassingly simple question: what am I protecting?

Fifteen years ago, that question had a reasonably tractable answer. Your IT environment was largely on-premises, mostly static, and bounded by a physical perimeter. Asset management was difficult, but it was at least a finite problem. You could enumerate your servers, your endpoints, your network devices, and your applications, and arrive at something that roughly represented reality.

The shift to cloud-native architectures has fundamentally transformed the nature of the attack surface. Infrastructure is now spun up and torn down continuously, often without the involvement of a central IT team. Developers, often using AI and third-party packages, deploy directly to production. SaaS adoption has distributed applications and data across dozens of third-party platforms. Corporate acquisitions bring inherited infrastructure that may never have been properly audited. Your attack surface now extends to every domain, IP, cloud service, API, exposed service, and third-party touchpoint that can be linked back to your organization.

The result is asset debt: accumulated exposure from infrastructure that exists and is discoverable by attackers, but that has never properly entered the organization's field of view. These assets are invisible to the security team, but they are not invisible to threat actors. Attackers do not respect your internal categorization of what is "in scope."

The challenge is compounded by the pace of modern engineering. In organizations that operate with a move-fast culture, where code ships multiple times per day and infrastructure is treated as ephemeral, the attack surface is changing faster than any periodic review process can track. A monthly scan that covers the assets you already know about is not even asset management, it's a historical record.

What "True" Attack Surface Management Actually Means

Given how the term has proliferated across the security market, it is worth being precise about what Attack Surface Management actually means. And, equally, what it does not mean.

ASM is not a tool. It is not a dashboard that aggregates data from other security products. It is not a scheduled scan of a predefined asset list. It is a practice: a continuous, systematic, and comprehensive approach to knowing your external attack surface as an attacker would see it. At its core, true ASM is built on four foundational principles: asset awareness and visibility, speed, signal quality, and scale.

Asset Awareness & Visibility

Means knowing what exists. Not just the assets you know about, but the full scope of your external footprint, including assets you did not deliberately create or may not even know you own.

Speed

Means operating in as close to real-time as possible. New assets are deployed constantly; new vulnerabilities are published daily; the window between a new exposure appearing and an attacker discovering it can be measured in hours.

Signal Quality

Means that the findings your ASM program surfaces are accurate, high-fidelity, and confirmed exploitable; not a volume of noisy alerts that security teams must spend hours validating before they can act.

Scale

Means that your ASM capability can cover the entire attack surface continuously, not just a representative sample.

The distinction between true ASM and the asset discovery tools that dominate much of the market is significant. Discovery is just the starting point. More important is everything that comes after: exposure monitoring, vulnerability analysis, prioritization, and remediation. An ASM program that stops at discovery has solved perhaps a quarter of the problem (really, it hasn't solved the ultimate problem at all). Critically, this also means that ASM is an ongoing practice, and one that requires engineering rigor to get right.

Many organizations have attempted to stitch together open-source tools for discovery, vulnerability scanning, and various feeds for threat data, only to find that the coordination between those components is itself a significant engineering challenge. Two years later, the project has consumed considerable resource and still has meaningful gaps. It's not to say that open-source tools are insufficient; but that the integration, reliability, and safety mechanisms required to run continuous, comprehensive ASM at scale are non-trivial to build and maintain.

The ASM market has made this harder by introducing significant definitional confusion. Buyers evaluating the category will encounter products that call themselves ASM but are primarily dashboards pulling from existing security tools, products that perform pure-play discovery without any exposure analysis, consulting-led services that perform a point-in-time assessment, and occasionally true end-to-end platforms that combine continuous discovery with deep vulnerability analysis. From the outside, these can appear superficially similar. In practice, they are solving fundamentally different problems. You might have evaluated multiple vendors in this category and nearly abandoned the space

entirely. Probably because nothing seemed to deliver what you needed. To find it, let's first outline the most fundamental parts of the attack surface layer of PTM.

Seeing Yourself Through an Attacker's Eyes

The most useful mental model for understanding what ASM is trying to achieve is to think about it from the attacker's perspective. A capable attacker doesn't start with your asset inventory. They start with what they know about you: a company name, a domain, an IP address, or another known identifier. From that initial seed, they conduct reconnaissance, expanding outward to discover domains and subdomains, IP ranges, cloud assets, exposed services, technologies in use, and data that has been leaked or inadvertently exposed. Each discovery can lead to another, allowing attackers to build an increasingly complete picture of your external exposure. Continuously, methodically, and at scale.

This is the inside-out view. True ASM attempts to replicate exactly this reconnaissance process, emulating what an attacker would find. The goal is to ensure that the security team sees their own attack surface before the attacker does, and that the resulting picture is as comprehensive and accurate as the attacker's would be.

What this reveals, invariably, is a set of exposures that existing security processes have missed. And critically, many of the most significant exposures on any given attack surface are not things that traditional vulnerability management would catch. One of the most important, and underappreciated limitations of conventional security programs is the overreliance on CVEs as the primary exposure signal.

A CVE is a standardized identifier assigned to a publicly known cybersecurity vulnerability. CVE-based vulnerability scanning is useful for identifying known vulnerabilities in software, but it creates a significant blind spot for exposures that do not fit that model.

Misconfigurations, an exposed admin panel with default credentials, or many other vulnerabilities exploited by threat actors can exist entirely outside of CVEs. The same limitation applies to newly discovered vulnerabilities. A zero-day can exist before it has been publicly disclosed, patched, or assigned a CVE. By definition, an organization cannot rely on a CVE identifier to identify an exposure that has not yet entered that public vulnerability ecosystem.

This is why, as we outlined in the first part of this book, offensive security and vulnerability research are so critical to ASM in the current era. Those with the best grasp of what ASM should be, tend to have roots in red teaming, penetration testing, and bug bounty work, precisely because those disciplines train practitioners to think like attackers, not like patch managers. The techniques for enumerating and identifying exposures come directly from offensive tradecraft: the same methods used to find bugs in external program are the methods applied to continuously monitor an organization's attack surface.

The “Management” Part of Attack Surface Management

Perhaps the most underappreciated dimension of ASM is the part that the name gestures at explicitly but that most implementations get wrong: the management.

Finding things is hard. But making findings consumable, trustworthy, and actionable is where most ASM programs, and most ASM products, fall short. When signal quality is poor, security teams that cannot trust their findings have to manually validate every result before escalating to engineering or operations. The tool designed to save time has created a new class of work. Worse, it has eroded confidence in the program itself, and a security team that has been burned by false positives will start to discount real findings alongside the noise.

If something is proved exploitable, however, that is the signal amongst the noise. Instead of triaging findings before they can be escalated, they go directly to the engineering team responsible for the asset with a clear, evidence-based remediation request. The cycle time from discovery to remediation compresses significantly.

It is also worth being honest about the limitations of purely technical scoring when it comes to prioritization. CVSS scores, for all their ubiquity, are a flawed prioritization mechanism. Research has shown that prioritizing vulnerabilities solely by CVSS score can be little more effective than random selection when it comes to reducing real-world exploitability³.

Because CVSS measures theoretical severity, not actual exploitation likelihood in a specific environment. True prioritization requires context: which of these exposed assets is most likely to be targeted? Which vulnerabilities are actively being exploited in the wild? Which exposures connect to your most sensitive systems? Those questions cannot be answered by CVSS alone.

“ *The real attack surface is alive and constantly expanding. It is almost always significantly larger than anyone expects. And it contains the assets that attackers will find first.* ”

³[1] Jacobs, J., Romanosky, S., Edwards, B., Adjerid, I. and Ablon, L., Improving vulnerability remediation through better exploit prediction, Journal of Cybersecurity, Vol. 6, Issue 1, 2020.

The Limits of ASM Alone

Attack Surface Management, even done well, has a fundamental limitation: it tells you what exists and where you are exposed, but it cannot tell you what attackers are actively planning to do with that information.

An ASM program will give you a comprehensive, continuously updated picture of your external attack surface. It will surface exposures that your existing security tools missed. It will give you a credible inside-out view of your own organization. But without additional context; intelligence about real-world attacker behavior, intent, and targeting, it is not yet possible to answer the question that matters most to a security team with finite resources: which of these should we fix first?

Data without context is not actionable intelligence. A thousand exposures ranked by CVSS score is just a work queue with no confident starting point. Deciding which exposures represent the most imminent, material risk to the organization requires visibility into the threat landscape that exists outside the organization's own infrastructure. It requires understanding which vulnerabilities are being actively weaponized by threat actors, and knowing whether your organization, your sector, or your technology stack is being specifically targeted. It requires intelligence about attacker activity that takes place outside your perimeter.


This is where the inside-out view of ASM meets its necessary counterpart: the outside-in view of threat intelligence. ASM tells you what your attack surface looks like from the outside. Threat intelligence tells you what threat actors are doing to attack surfaces like yours. Together, they form the basis for a genuinely prioritized, context-driven exposure management program. The security team is not just aware of what's exposed, but understands which ones attackers are most likely to exploit, and can act preemptively on that basis. More on that to come...

ASM as the Foundation of PTEM

Every pillar in the Preemptive Threat Exposure Management framework depends, ultimately, on knowing what you own. Threat intelligence, however sophisticated, needs to be relevant to your organization's attack surface. Exposure prioritization is meaningless if the most exposed asset in your estate is one you can't see. Remediation workflows presuppose that there is something to remediate: a specific, identified asset with a specific, confirmed exposure.

ASM is the common data layer on which all other PTEM capabilities are built. It is the mechanism by which the organization achieves the foundational honest answer to the question: what are we actually trying to protect, and what does it look like from the outside?

The maturity journey toward preemptive security starts here. Organizations typically begin with reactive asset discovery: acknowledging that their existing inventory is incomplete and beginning to close the gap. The next stage is continuous ASM: moving from point-in-time discovery to an ongoing program that tracks the attack surface as it changes in real time. The final stage, and the one that enables PTM, is ASM integrated with threat context, a state in which the complete, continuously updated picture of the attack surface is mapped against real-world attacker intelligence to produce a prioritized, evidence-based view of where preemptive action is most needed.



That final stage is what the rest of this book is about. But it starts here, with the discipline of really knowing what you own, and what's exposed.

Attacker Reality: The Threat Intelligence Layer

Most organizations consume threat intelligence. But having access to information about threats is not the same thing as having intelligence that can inform action.

Security teams receive indicators of compromise, reports on threat actors, vulnerability advisories, intelligence on emerging campaigns, and warnings about activity across the threat landscape. All of these can be valuable. The challenge is understanding what they mean for a specific organization, its attack surface, and the decisions it needs to make.

That distinction matters because intelligence is more than information about what is happening in the threat landscape. The importance comes in understanding what that information means in a particular context, and what should be done as a result. A piece of information can be accurate, timely and technically interesting without necessarily being actionable intelligence for the organization consuming it.

This is one of the reasons threat intelligence can sometimes feel disconnected from day-to-day security operations. Organizations may have access to an enormous amount of information about threat actors, vulnerabilities and malicious infrastructure, but still struggle to answer the questions that matter most: what does this mean for us, how urgent is it, and what should we do about it?

The problem is not that threat intelligence is a bad idea. The concept is vital, and the premise is correct. The challenge is turning information into understanding, and understanding into action. That requires context: knowing what the information means in relation to your organization, your technology, your exposure, and the threats that are most relevant to you.

This chapter is about where information about attacker activity comes from, how it becomes intelligence, what makes that intelligence actionable, and why it becomes particularly powerful when it is applied to the specific, known attack surface we explored in the previous chapter. It introduces the second foundational pillar of Preemptive Threat Exposure Management: the attacker reality layer.

Information Isn't Intelligence

Not everything described as threat intelligence is intelligence in the same sense.

Data is the raw observation: an IP address, a domain, a leaked credential, a vulnerability, a forum post, a threat actor's name. Information adds context to that observation: when it was seen, where it came from, how it has been used, who is associated with it, or what activity surrounds it. Intelligence goes another step further. It involves understanding what that information means in a particular context, assessing its significance, and using that understanding to inform a decision.

The distinction is key. An IP address associated with malicious activity may be useful information.

Knowing that the infrastructure belongs to an attacker who is currently targeting organizations in your sector is more meaningful. Knowing that the same attacker is actively seeking access to a technology that you have exposed to the internet begins to create a very different picture. If you can connect that activity to a specific exposed asset in your own environment and determine what action should be taken, the information has become actionable intelligence.

This is why the quality of threat intelligence cannot be judged simply by the volume of information produced. A thousand indicators, reports or alerts are not necessarily more valuable than ten if those ten enable better decisions.

“ *Not everything described as threat intelligence is intelligence in the same sense.*

What are Attackers up to?

It is worth being precise about where attacker activity actually takes place: the sources that, taken together, constitute the attacker reality.

The picture is broader than many organizations realize. The visible web, the indexed internet that most people interact with daily, is a small fraction of the total online environment. Beneath it lies the deep web: content that is not indexed but is accessible through standard browsers, including databases, private forums, and password-protected sites. Below that (not literally, but it helps with the analogy) lies the dark web: a set of networks, most notably Tor, that require specialist software to access and that are deliberately designed to resist indexing, surveillance, and attribution. These environments collectively represent the majority of online activity that defenders cannot see and that attackers exploit precisely for that reason.


On the dark web, cybercriminal activity operates as a sophisticated, structured economy. Hacking forums serve as marketplaces and coordination hubs where threat actors discuss targets, share techniques, trade exploits, and recruit affiliates. Ransomware groups maintain leak sites where they publish the data of victims who have refused to pay, and operate backend infrastructure that functions more like a business than a criminal gang, complete with help desks, affiliate program, and reputation systems. Initial access brokers, a relatively recent and consequential development in the the grand scheme of things, sell verified footholds in compromised organizations to other criminal actors, creating

a division of labor in which the entity that breaches an organization is often not the same as the one that ultimately deploys ransomware or exfiltrates data.

Encrypted messaging platforms, particularly Telegram, have become an increasingly significant layer of the threat landscape over recent years. As law enforcement has become more effective at disrupting dark web forums, criminal coordination has shifted toward platforms that offer large group functionality, pseudonymity, and encryption. Ransomware groups, hacktivist collectives, and fraud networks now maintain active Telegram channels through which they announce attacks, share tools, recruit members, and communicate with affiliates. The intelligence value of this layer is substantial and is frequently underestimated by organizations that equate dark web monitoring with Tor-only visibility.

Beyond these layers, there is significant intelligence value in what might loosely be called the grey web: clear web forums where stolen data is advertised, paste sites where credentials and sensitive data are dumped, public code repositories where developers inadvertently commit sensitive configuration data, and the social media presence of threat actors who are often less operationally careful on mainstream platforms than they are on protected forums.

Each of these sources represents a layer of the attacker reality. Together, they constitute the environment in which attacks are planned, resources are acquired, and targets are selected, all before any malicious activity appears on the victim's own infrastructure.



The Dark Web: The Biggest Blind Spot

Of all the sources that make up the attacker reality, the dark web represents the most significant intelligence gap for most organizations. It's not the only source that matters, but it is the one that is hardest to access, hardest to interpret, and therefore most consistently absent from organizational threat visibility programs.

The barriers are real. Dark web sources are not indexed, do not respond to conventional web crawlers, and require specialist tooling and infrastructure to access securely. Manual monitoring is not a viable approach at organizational scale: it is time-consuming, operationally complex, and carries meaningful risk for analysts who operate without proper experience, controls, and tradecraft. Forums appear and disappear. Threat actors move between platforms. Content is ephemeral, often deliberately so. Languages and dialects used across criminal communities range from Russian and Chinese to coded slang that requires contextual expertise to interpret correctly. And the volume of content across even a narrow slice of the dark web is far beyond what any manual process could meaningfully analyse.

The practical consequence is that most organizations have no visibility into what is happening in the environment where attacks are staged. This means they are consistently late. A few examples of this would be:

- They discover that credentials were stolen when those credentials are used to access their systems, not when they were listed for sale on a dark web marketplace.
- They discover that a vulnerability in their stack is being actively exploited in the wild when their own systems are affected, not when that vulnerability was first being traded or discussed in criminal forums.
- They discover that they are a named target of a ransomware group when the group deploys ransomware, not when the early warning signs appeared weeks earlier.

Shifting Left in the Cyber Kill Chain

The Cyber Kill Chain, developed by Lockheed Martin, and the MITRE ATT&CK framework both describe the stages through which a cyberattack progresses from initial intent to impact. The earliest stages, reconnaissance and resource development, are the phases in which an attacker gathers information about their target, acquires tools and access, and prepares the infrastructure for the attack. These stages take place entirely outside the victim's network. No matter how good it is, no security tool that monitors internal infrastructure has any visibility into them.

Dark web monitoring sits at this crucial layer, in the earliest stages of an attack forming. When a threat actor conducts reconnaissance on a target organization using dark web sources, that activity generates a signal.

The role of threat intelligence is to interpret those signals: to establish what they reveal about attacker capability, intent or opportunity, how credible the information is, and whether it has a meaningful connection to the organization being defended.

This is where the attacker reality layer becomes more than simply another source of threat data. Its value lies in making the connection between what attackers are doing and what that activity means for the organization. The earlier that connection can be made, the more options the defender has.

That temporal advantage is the entire value proposition of the attacker reality layer in PTEM. It is also what fundamentally distinguishes it from every reactive security tool in the traditional stack. SIEM, XDR, and network detection and response all have enormous value as part of a defence-in-depth architecture. But they are all designed to detect activity that is already occurring within the network. They are built to detect when something is happening, but to be preemptive, shouldn't we be asking whether something is about to happen?

The security principle at work here is sometimes described as “shifting left”: moving the point of detection and response earlier in the attack lifecycle, reducing the window in which an attacker has freedom to operate and increasing the time available to prepare or prevent. A credential exposure detected on a dark web marketplace before the credentials are used is an exploitable condition that can be closed. The same exposure detected after a breach requires incident response, forensic investigation, regulatory notification, and reputational management. Hardly equivalent problems.

Intelligence That Hits Home

Understanding the threat landscape in the abstract is valuable for building general defensive strategy; if you don't look outside your own four walls at all, you'll soon lose perspective. But for the purposes of defending an organization against threats day-to-day, intelligence needs to be specific. Intelligence that is explicitly relevant to your organization is far more actionable than information about the general threat environment. A threat feed that tells you ransomware attacks on financial services organizations increased by forty percent last year is useful context. An alert that a known ransomware affiliate has posted a thread on a dark web forum discussing organizations with your technology profile, in your sector, is much more significant.

Organization-specific monitoring is one way of achieving this specificity. By monitoring the wider attacker environment for mentions of your organization's name, domains and IP ranges, technology stack, executive names, supplier ecosystem and sector, it becomes possible to filter the enormous volume of dark web and threat data down to the subset that is directly relevant to your risk. The signal only becomes intelligence when its significance to your organization is understood.

ASM and Threat Intelligence: Greater Than the Sum of Their Parts

The previous chapter established that Attack Surface Management, even done well, is limited when conducted in isolation: it tells you what exists and where you are exposed, but it cannot tell you what attackers are actively planning to do with that information. The same is true in reverse for threat intelligence.

Threat intelligence can tell you what attackers are doing, what they are interested in, and what capabilities or intentions they may have. But without an accurate and continuously updated understanding of your own external attack surface, it can be difficult to determine what that information means for you. The intelligence may be real, credible and important, but the next course of action can still be unclear.

Consider an example. An Initial Access Broker lists verified access to an organization through an exposed RDP service. The threat intelligence is real. The seller has authenticated access and is offering it to other criminals. But without knowing which asset the access relates to, whether it is still exposed, what sits behind it, and what other exposures might be connected to it, the security team has limited ability to assess the risk or respond.

This is the core argument for PTEM as a comprehensive framework rather than a collection of individual tools: ASM provides the organizational context to which threat information can be applied. Threat intelligence provides the external

context that helps explain which exposures matter most. Together, they turn exposure data and attacker signals into a prioritized understanding of risk.

When those two views are continuously connected, the organization can move beyond simply knowing what is exposed or what attackers are doing. It can understand where those two realities intersect.

What Actionable Looks Like

The ultimate test of any intelligence program is the quality of the decisions it enables. In the context of PTM, the attacker reality layer is valuable because it can turn pre-attack signals into an understanding of what matters and, ultimately, into preemptive action. For example:

-  When an employee's credentials appear on a dark web marketplace, the action is clear: identify the affected account, rotate credentials immediately, audit recent access logs for any signs of prior use, and assess whether those credentials provided access to sensitive systems that require further investigation. The organization acts before the attacker does.
-  When an exploit for a vulnerability in a technology the organization runs is being actively traded or discussed in dark web forums, the action is also clear: escalate the relevant finding from the ASM program to the top of the remediation queue, regardless of its theoretical risk, because attacker interest is a real-world demand signal that CVSS does not capture. Patch or mitigate before exploitation.

3

When a ransomware group's forum activity suggests they are actively selecting targets in a specific sector, a security team with visibility into that activity can build or update a threat model specific to that group, map their known tactics, techniques, and procedures against their own defensive controls, and address the specific gaps that make them a viable target. The attack has not happened and the opportunity to prevent it is still open.

4


When dark web or Telegram signals indicate that an insider may be exploring or advertising access to the organization's systems, the security team has the opportunity to conduct an investigation, identify the individual, and act before any data leaves the organization or any access is sold. The alternative is discovering the insider threat during a forensic investigation after the fact.

In each case, the organization has gained back time. That time is the difference between a preemptive security posture and a reactive one.

The Foundation Is In Place

With the inside-out view of ASM and the outside-in view of threat intelligence working in concert, the core pillars of PTEM are in place. The organization knows what it owns and what it looks like from the outside. It knows what threat actors know about it and what they are planning. It has a prioritized, evidence-based picture of which exposures represent the most imminent, material risk.

What it needs now is a systematic, operationalized way to act on that picture. The rest of this book shows how exposure and attacker intelligence can be turned into preemptive action in practice, at the pace and scale that modern threat environments demand.



Ransomware — Where Preemptive Defense Matters Most

The term “ransomware attack” is something that, over the past 10 years, has become commonly understood, and conjures up a specific image: screens locked, systems offline, a ransom note demanding payment in cryptocurrency. For the security team scrambling to respond, it feels like the attack has just begun. In reality, it ended there. Everything before that moment, the scanning, the access, the persistence, the data exfiltration, that was the attack. The point that files were encrypted was simply the bill arriving.

This distinction matters enormously. It means that virtually every aspect of a conventional ransomware response, such as the incident response plans, the backup restoration procedures, the negotiation playbooks, is optimized for the wrong moment. By the time any of those processes activate, the ransomware group has already done what it came to do.

Ransomware is the threat type that makes the case for preemptive security most viscerally. Not because it is the most technically sophisticated attack, (often it isn't), but because the gap between where defenders focus their attention and where attackers actually operate is particularly wide.

A Ransomware Attack in Stages

Understanding why preemptive defence is so effective against ransomware requires understanding what ransomware operators actually do, and when they do it.

Modern ransomware is not a single event. It is the final stage of a structured, multi-phase operation that more closely resembles a business process than a hack. An example could look like this:



Reconnaissance

Attackers may begin by scanning the internet for exposures: open RDP ports, vulnerable VPN appliances, misconfigured cloud storage, forgotten subdomains, legacy systems left accessible but unmonitored. Much of this reconnaissance can be automated, and continues as attackers map systems, accounts, services, and valuable data.



Initial Access

Entry can be gained through one of a predictable set of vectors: compromised credentials (often acquired through infostealers or purchased from initial access brokers operating on dark web forums), exploitation of unpatched vulnerabilities, or exposed remote access services. Sophisticated zero-days are sometimes used, but ransomware operations also routinely succeed through weaknesses that are visible, and fixable, from the outside.



Persistence and privilege escalation

Once inside, attackers embed themselves quietly, escalating privileges and establishing footholds across multiple systems. This phase can last days or weeks. Dwell time is a significant advantage.



Lateral Movement

Critical data stores, domain controllers, and backup systems are identified and mapped. The attacker understands your environment, including privileged accounts, critical systems and recovery mechanisms.



Data Exfiltration

Sensitive data is staged and stolen. This enables double or triple extortion: the ability to threaten publication of data even if an organization recovers from encryption. It is why backup-and-restore strategies, while necessary, are no longer sufficient on their own.



Encryption & Extortion

Only now does the visible attack begin. By this point, the outcome is largely determined. Sure, you can get your systems back up and running, but having exfiltrated your sensitive data, the attacker still holds the cards.

The ransomware ecosystem that delivers this process has matured considerably. Ransomware-as-a-Service (RaaS) models now dominate among the most prolific groups, with professional affiliate programs, defined operating standards, and specialized roles. Initial access brokers operate as a distinct market, selling verified footholds to ransomware operators. What was once an opportunistic criminal enterprise has industrialized. Our own threat intelligence team's research into the ransomware landscape has shown victim numbers growing year on year, groups proliferating despite law enforcement pressure, and the barrier to entry falling as automation and AI enable newer actors to scale operations rapidly. The threat is not slowing down.

Where Defenses Break Down

The previous chapters established why reactive security fails as a general model. Ransomware illustrates that failure in concrete terms.

Traditional security programs are oriented inward. They monitor network traffic, trigger alerts on malware execution, and respond to anomalies detected inside the perimeter. The problem is that the earliest and most disruption-amenable phases of a ransomware attack happen entirely outside that perimeter. They happen in the attacker's reconnaissance of your external footprint, in dark web markets where your credentials are circulating, or in the gap between a vulnerability being disclosed and your team patching it.

By the time an internal alert fires, the attacker has already passed through the stages where intervention would have been cheapest and most effective.

There is also a visibility problem that compounds this. Internal asset inventories rarely reflect the true external attack surface. As you've now read several times, shadow exposure, cloud deployments that expand faster than security teams are notified, third-party integrations, and legacy infrastructure all create exposure that organizations don't know they have. Ransomware operators, scanning continuously at scale, find them readily. If you cannot see what an attacker sees, you cannot meaningfully reduce the risk they represent.

Preempting Ransomware: Moving Left of the Attack

Applying these capabilities to ransomware reveals a set of specific, high-value intervention points that map directly onto how ransomware groups operate.

Closing the reconnaissance window

Ransomware operators can gain an advantage during reconnaissance when the assets they discover are unknown, misclassified, or insufficiently monitored by the organizations that own them. Continuous attack surface monitoring reduces that gap. When an exposed RDP port or vulnerable appliance discovered from the outside also appears in your own external view, you have an opportunity to remediate it before an attacker turns it into an entry point. The intervention is straightforward, acting on visibility you may have previously lacked.

Cutting off initial access

Compromised credentials are an important entry point for ransomware operations. Dark web monitoring can provide an early warning when employee credentials surface in criminal marketplaces, forums, or infostealer logs, potentially before those credentials are used against the organization. That window, however narrow, creates an opportunity: invalidate the credentials, investigate the affected account, and prevent them from becoming an initial access point. For ransomware operations that rely on stolen credentials, this is an opportunity to interrupt the attack chain before access is established.

Knowing what ransomware groups are actually exploiting

Threat intelligence on active ransomware groups, the access methods they are exploiting, and who they are targeting can transform how security priorities are set. A vulnerability that might otherwise sit in a backlog can become an urgent fix when attacker behavior indicates it is being actively targeted. This is what distinguishes a preemptive posture from simply running a vulnerability management program: the outside-in, threat-informed perspective that narrows attention to what actually matters right now.

The Practical Difference

Consider two organizations facing the same ransomware threat actor. Both have broadly similar technology environments.

The first organization maintains a reactive security posture. Its asset inventory was last reviewed months ago. Credentials belonging to several employees, compromised by an infostealer several weeks prior, are circulating on a dark web forum. One of its VPN appliances is running a version with a known vulnerability that the relevant ransomware group has been actively exploiting. None of this is visible to the security team. The first indication of a problem is the point of encryption.

The second organization has continuous external attack surface monitoring in place, integrated with dark web credential monitoring and ransomware threat intelligence. The compromised credentials surface in monitoring within hours, if not instantly.

They are invalidated and the affected accounts investigated. The vulnerable VPN appliance is flagged as a priority not because of its CVSS score in isolation, but because threat intelligence confirms active exploitation by the group most likely to target organizations in their sector. It is patched before any scanning activity reaches it.

The ransomware group finds no viable entry point and moves on. This brings home what preemptive cybersecurity is really all about in simple terms: disappointing attackers at every step.

The tale of these two hypothetical organizations shows the importance of shifting your focus and investment to the stages of the attack chain where intervention actually makes the strongest difference.

**“ *Preemptive cybersecurity
is about disappointing
attackers at every step* ”**


Ransomware Is Not Inevitable



For many organizations, ransomware has become a threat to be prepared for rather than a threat they expect to prevent. Faced with an evolving and increasingly difficult threat landscape, security teams and boards have understandably invested heavily in resilience: reliable backups, incident response plans, cyber insurance, and ransom negotiation retainers. These measures are beneficial, but their prominence can reflect a pragmatic assumption that prevention is impossible.

Law enforcement operations against ransomware groups continue to show success, and prove that putting continued pressure on these groups works, but as we see from the data and the record numbers of victims, it cannot be the only solution. Ransomware will continue to evolve. Groups will rebrand, techniques will adapt, and new vulnerabilities will emerge.

Resilience capabilities remain important. The ability to recover quickly from a successful attack is a legitimate part of a mature security posture. But recovery and preemption are not an either/or choice, and treating resilience as the primary strategy, while underinvesting in preemption, accepts a level of damage and disruption that is avoidable. Ransomware attacks carry consequences: operational downtime, regulatory exposure, reputational harm, extortion risk, that no backup strategy fully mitigates.

Focus on intervening earlier, at the stages of the attack chain where the cost of action is low and the potential impact is high. The message on ransomware is that you do not have to resign yourself to fighting it after the fact.





| Building a Preemptive Organization

Understanding PTEM Is Not the Same as Implementing It

Throughout this book we've explored why cybersecurity needs to change. We've examined the realities of today's threat landscape, looked at the world through the attacker's eyes, and introduced Preemptive Threat Exposure Management (PTEM) as a practical framework for finding and reducing the exposures that matter most.

By now, you should understand what preemptive cybersecurity is and why it represents a more effective way of defending modern organizations.

But how one implements these new practices requires further thought.

Organizations invest in Attack Surface Management. They consume threat intelligence. They deploy new exposure management platforms and automate vulnerability scanning. Yet despite all of this, their security operations remain fundamentally reactive. They still prioritize work based on static severity scores, measure success by the number of vulnerabilities closed, and discover critical exposures only after an attacker has already found them.

The technology has moved on, but the organization hasn't.

This is because being preemptive is not simply about trying another security program or another technology category. It changes how an organization identifies risk, makes decisions, assigns ownership and measures success.

The preemptive cybersecurity model that PTEM enables isn't achieved by something you install but something your organization becomes.

This chapter explores how to build that kind of organization.

“

*The
technology
has moved
on, but the
organization
hasn't.*

PTEM Is an Operating Model, Not a Project

Many cybersecurity initiatives fail because they are treated as projects. Projects have a beginning and an end. They introduce new tooling, update a process or satisfy a regulatory requirement before attention shifts elsewhere.

PTEM is different in that it changes how security decisions are made every day.

Traditional vulnerability management is largely transactional. Assets are scanned. Vulnerabilities are identified. Tickets are generated. Teams work through remediation backlogs until the next scan starts the cycle again.

Rather than asking “What vulnerabilities do we have?”, your organization begins asking: “Which exposures are most likely to be exploited, and how quickly can we resolve them?”

Instead of measuring activity, organizations begin measuring outcomes, and continually search for risk rather than waiting for it to rear its head. The most sophisticated PTEM platform in the world cannot compensate for an organization that still thinks reactively.

Five Shifts That Define a Preemptive Organization

Becoming preemptive doesn't require replacing every existing security process overnight. Instead, it requires changing a handful of fundamental behaviors that influence thousands of daily decisions.



Shift One

From Vulnerability Management to Exposure Reduction

A vulnerability is not the full story in isolation. What matters is whether it contributes to an exploitable path towards something valuable.

This changes both prioritization and behavior. Instead of asking how many vulnerabilities remain open, organizations ask which exposures create the greatest opportunity for attackers.

The objective goes from reducing vulnerability counts to reducing real attack paths.



Shift Two

From Security Ownership to Shared Ownership

Funnily enough, it isn't the security team that owns security. Stay with us on this one...

Security teams own visibility. They provide expertise. They identify exposures. But they rarely wholly own the infrastructure, applications or identities that ultimately require remediation. A preemptive organization recognizes that exposure management is inherently cross-functional.

- Infrastructure teams reduce operating system risk.
- Cloud teams reduce configuration risk.
- Engineering teams reduce application risk.
- Identity teams reduce privilege risk.
- Security orchestrates these activities, but it cannot perform them alone.

Ownership therefore shifts from being departmental to organizational. This helps security become an enabler rather than a bottleneck.

→ Shift Three

From Periodic Assessment to Continuous Awareness

Attackers don't wait for the next weekly or quarterly vulnerability scan, so why are you? Every cloud deployment, software release, third-party integration and identity change alters an organization's exposure. PTEM therefore replaces periodic visibility with continuous awareness. Higher frequency scanning is great, but more broadly, this shift builds an organization that expects change, monitors change and responds to change as part of normal operations.

Visibility becomes continuous because the attack surface is continuously changing.

→ Shift Four

From Compliance to Risk Reduction

Compliance remains important. Organizations cannot ignore regulatory obligations, industry standards or customer requirements. But compliance should be treated as a consequence of good security rather than its purpose. Reactive organizations think in terms of passing audits, but preemptive organizations are focused on denying attackers.

Interestingly, organizations that consistently reduce real exposure often discover that compliance becomes significantly easier.

They have

- Fewer unmanaged assets.
- Better visibility.
- Stronger governance.
- Faster remediation.

Compliance becomes evidence that security is working, not the reason it exists.

→ Shift Five

From Reacting Quickly to Thinking Earlier

As we outlined right at the start of this book, much of modern cybersecurity celebrates speed. Speed is great, it means faster detection, faster response, and faster recovery.

These capabilities remain essential, but they all begin after an attacker has already started their work.

If the question you're asking is how to prevent the attacker from succeeding in the first place, the ultimate measure of success is how many incidents never occur because the opportunity was removed before exploitation became possible.

By this point in the book this probably seems obvious, but depending on the organization and its maturity, it could be the biggest cultural shift of all.

Success is measured not by how well you respond to attacks, but by how effectively you deny attackers opportunities altogether.

What Changes Across the Organization?

A preemptive organization needs both technology and people working together.

That means PTEM cannot remain confined to the security team. Every function that influences organizational exposure has a role to play.

Leadership must redefine success: Rather than asking how many vulnerabilities remain open or whether patching SLAs have been met, executives should ask where the organization's greatest exposure lies and whether it is being reduced over time. They shape priorities, allocate investment and remove organizational barriers. Without leadership support, PTEM quickly becomes another security initiative competing for attention.

Security teams transition from acting primarily as vulnerability managers to becoming coordinators of organizational risk reduction. They expand beyond managing vulnerabilities to understanding attacker behavior, focusing in on the exposures that matter most and enabling the rest of the organization to act.

Engineering teams begin considering exposure as part of software delivery rather than after deployment. Cloud teams continuously manage configuration drift. Infrastructure teams prioritize remediation according to exploitability rather than arbitrary patch cycles.

Every team understands how its own decisions influence attacker opportunity.

Preemptive cybersecurity strategies succeed when exposure management becomes everyone's responsibility, even if accountability remains clearly defined.

“ *Every function that influences organizational exposure has a role to play.*

Building the Operating Model

Successful organizations reinforce cultural change with operating structures that make preemptive behavior the easiest behavior. That means establishing clear ownership for different exposure types, integrating PTM into existing operational workflows and creating governance that removes ambiguity rather than introducing bureaucracy.

- Security should own visibility, context and validation, empowered by preemptive capabilities that make this process rapid, relevant and high-signal.
- Operational teams should own remediation within their domains, acting with confidence due to the prior prioritization and context delivered by security
- Security and IT Leadership should own overarching prioritization, investment and accountability, having gained board buy-in based on material evidence of security outcomes and a clear path to steering future investments.

These responsibilities should already exist within mature organizations. PTM simply aligns them around reducing attacker opportunity instead of managing vulnerability volume. Technology should support these workflows, not dictate them.

Exposure data should flow naturally into IT service management platforms, engineering backlogs and change management processes. Validation should occur automatically wherever possible. Threat intelligence should continuously refine prioritization. Automation should remove repetitive work without removing human judgement.

When PTEM fits naturally into existing operational processes, it stops feeling like another security program and starts becoming part of how the organization operates.

Culture Follows Incentives

Organizations become shaped by what they reward. If engineering teams are measured heavily on release velocity, security can become perceived as friction when it slows delivery. If security teams are rewarded primarily for generating and closing tickets, activity can become a more visible measure of success than the risk those activities actually reduce.

Building a preemptive organization therefore requires changing incentives as much as changing technology. Celebrate teams that reduce meaningful exposure. Recognize collaboration across departments. Measure reductions in attacker opportunity rather than increases in security activity.

Instead of asking whether tickets have been closed, ask whether the underlying exposure has been removed. Instead of reporting vulnerability counts, report exposure reduction. Language shapes behavior, and behavior shapes culture.

Becoming Preemptive Is a Journey

No organization becomes fully preemptive overnight. A successful transformation can just begin with a single objective. That could be protecting internet-facing infrastructure, reducing ransomware exposure, or even just securing one business-critical application.

Early success builds confidence, which impacts organizational support, and in turn, wider adoption.

Over time, PTEM becomes less about introducing new technology and more about changing expectations. Teams naturally begin thinking about attacker behavior during planning meetings, software releases and infrastructure changes. Exposure management becomes part of everyday decision-making rather than a separate security activity.

Building a preemptive organization requires changing how decisions are made across the business. The people, processes and leadership are what determine whether the capabilities of PTEM translate into meaningful reductions in risk.

Once an organization begins thinking differently, it also needs to measure success differently. Traditional metrics such as vulnerability counts, patch compliance, and ticket volumes tell us how busy security teams have been, but not whether the organization has become more resilient.

The next chapter explores what organizations should measure instead, and how the right metrics reinforce the behaviors that make preemptive cybersecurity successful.

Why Today's Cybersecurity Success Metrics Miss the Point

There is a particular kind of confidence that comes from a green dashboard. The numbers, on their face, look good. And yet, somewhere in the background, all it takes is one unknown exposure to undo everything.

This is the central problem with how many organizations measure success. Not that they measure badly, but that they often measure the work being performed rather than the risk being reduced.

Busy Isn't the Same as Effective

The metrics that dominate security reporting are often operational metrics. Ticket volume. Time to close. Mean time to detect. Mean time to respond. Vulnerabilities remediated.

These metrics tell you how much work a security function is doing and how quickly it is doing it. They do not necessarily tell you whether the business is actually safer.

The damage runs deeper than bad reporting. When analysts are measured on ticket throughput and closure speed, their behavior changes to match. Subtle, low-confidence alerts get dismissed because sitting with uncertainty costs time. Deep-dive investigations get deprioritized because they don't move the numbers. Exposure patterns that emerge slowly across weeks of signals get missed because no single ticket captures the whole picture.

The metric shapes the work. When the metric is wrong, so is the work it produces.

“

The metric shapes the work. When the metric is wrong, so is the work it produces.

Even Outcome-Focused Metrics Can Come Too Late

Many security leaders have arrived, rightly, at the conclusion that most security metrics are noise. Their answer has been to strip everything back to the essentials: does the security organization detect attacks quickly, and does it respond when it does? These fundamentals; time to detect, time to respond, percentage of incidents caught within a defined window, are at least trying to capture something real.

But these metrics generally begin their measurement once suspicious activity has entered the detection and response process. By the time your SOC detects an intrusion, an attacker may already have gained access or begun operating inside your environment. The SOC's response may be fast. It may even be excellent. But detection and response measure how effectively you deal with a threat once it has reached the point of detection, rather than whether the conditions that enabled the attack existed in the first place.

Timely detection is, in part, a damage-limitation metric: it measures how quickly you can identify and contain activity that has already reached your detection environment. That is a meaningful thing to be good at, but it should not be the only thing at the top of your security scorecard.

What Gets Rewarded Gets Done

Walking through some common security KPIs reveals the issue:

Tickets processed and closed rewards the person who clears the most tickets, fastest, and can quietly penalize the one who spends an afternoon tracing a complex exposure that turns out to be the organization's most significant risk that month. Over time, this metric can produce teams that are excellent at clearing queues and poor at finding what matters.

Vulnerabilities closed can create a similar incentive. A team can make impressive progress against a vulnerability backlog without necessarily reducing meaningful risk. Closing hundreds of low-risk findings may look better on a dashboard than spending a week investigating a small number of critical, externally exploitable vulnerabilities, even when the latter work does far more to reduce the organization's actual exposure.

Assets discovered can create another misleading sense of progress. Increasing the number of assets in an inventory may demonstrate improved visibility, but it does not necessarily mean the organization has reduced its exposure. Discovering an unknown asset is valuable; understanding whether it creates an attacker opportunity and then reducing that opportunity is the outcome that matters.

None of these metrics, on their own, tell you whether the organization is safer. They are measuring whether the team is performing more of the same activity, faster.

Measuring Prevention Requires a Different Starting Point

Preemptive cybersecurity insists on measuring earlier in the threat lifecycle, at the point where attackers are still looking for a way in rather than already moving through one.

This means shifting the scorecard to questions like: How many exploitable exposures did we identify and close this month, before any attacker could act on them? How early in the attack chain did we intervene? Is our external-facing risk profile getting smaller over time, or larger?

The metrics that answer these questions look different to the ones most security teams report today. They will vary by organization and to whom they are being reported, but the selection below give you an overview of the types of metrics that can be tracked to show a measurable strengthening of security posture:

Mean Time to Remediation (MTTR)

This tracks the average time between an exposure being identified and it being remediated or otherwise neutralized. Unlike traditional “mean time to patch” metrics, MTTR can be scoped specifically to exposures that have been validated as exploitable, rather than the full raw backlog. Measuring remediation speed against everything, including low-risk findings that are unlikely to be exploited, can dilute the signal.

Exposure Window Trend

A single MTTR number is a snapshot. What matters more is the trend over time – is the average window shrinking quarter over quarter, or holding steady while the volume of exposures grows? Given that industry data shows Time to Exploit dropping into hours rather than days, a flat or worsening trend is a strong warning sign, even if absolute numbers look reasonable in isolation.

Exposure Reduction Rate

This tracks the net change in validated exploitable exposure over a defined period, such as monthly or quarterly. It provides a direct measure of whether the PTEM program is reducing the organization's exploitable exposure over time.

Attack Surface Growth vs. Exposure Growth

As organizations adopt new technologies, their attack surface naturally expands. The useful comparison isn't attack surface size in isolation, but whether exploitable exposure is growing in proportion to attack surface growth, more slowly, or ideally shrinking despite it. A PTEM program should aim to ensure that exploitable exposure does not grow in step with attack surface, and ideally declines even as the attack surface expands.

False Positive / Low-Value Remediation Rate

The percentage of remediation actions that, following validation, are found to have addressed exposures that were not practically exploitable or did not present meaningful attack paths. A shrinking rate here can indicate that validation and attacker-reality signals are being used more effectively upstream, before remediation work begins.

Exposure Lead Time

An estimate of the lead time between identifying and reducing an exposure and the point at which the vulnerability was publicly disclosed or actively exploited.

“ *When you replace activity metrics with outcome metrics, [...] You create space for the kind of deep investigation that surfaces real risk*

The Conversation the Board Actually Wants to Have

Part of the reason activity-based metrics persist is that they feel reportable. Numbers go up. Progress feels visible. But any board member worth the role should be unmoved by ticket volumes and closure rates. What they need to know is whether the business is more or less exposed to the threats that could actually damage it.

A preemptive cybersecurity metric framework gives CISOs a far more compelling story to tell.

Compare these two ways of reporting the same month's activity:

The old version: "We processed 11,400 alerts this month. Mean time to detect was 19 minutes. Mean time to respond was 34 minutes. We closed 98% of tickets within SLA."

The outcome version: "We identified and remediated 180 validated exploitable exposures this month, including a set of leaked employee credentials that had appeared on a dark web forum and a misconfigured partner access point that mapped directly to a critical internal system. Our validated exploitable exposure is 14% lower than it was at the same point last year."

One of these is a description of activity, and the other is a demonstration of value. Only one of them answers the question a board is actually asking: have we reduced our risk?

What This Means for the Practitioners in the Room

Most people who enter security do so because they want to make a difference. They are drawn to the complexity of attacker behavior, the challenge of connecting signals that others miss, the intellectual satisfaction of getting ahead of a risk before it becomes an incident. All too often what they find is a productivity target and a queue.

When you replace activity metrics with outcome metrics, you do more than improve your reporting. You change what good looks like for the analysts doing the work. You create space for the kind of deep investigation that surfaces real risk. You reward the analyst who spends two days tracing an obscure exposure and closes it cleanly, rather than the one who clears the most tickets before end of shift. You build a culture around protecting the business, which is why most of these people joined security in the first place.

Preemptive cybersecurity is a more outcome-focused posture as well as a more effective one. It measures what actually matters: risk reduced, attacks prevented, exposures closed.

The Preemptive Cybersecurity Maturity Model

Understanding preemptive cybersecurity in principle is one thing; knowing where your organization actually stands is another. Most security programs don't move from reactive to preemptive overnight, they progress through distinct stages, often at different speeds across different capabilities. The maturity model that follows breaks this journey into five stages, from purely reactive security through compliance-driven programs, continuous exposure management, predictive validation, and finally to true preemption.

Each stage is assessed across ten dimensions:

1. Discovery
2. Validation
3. Prioritization
4. Attacker awareness
5. Speed of action
6. Data sources and signals
7. Remediation ownership
8. Metrics and success measures
9. Scope and coverage
10. Organizational integration

This will give you a practical framework to benchmark your current progress and identify the specific capabilities needed to progress. Rather than treating maturity as a single score, this model recognizes that organizations are rarely at one uniform stage; the goal is to build an honest picture of where you stand today, and a clear path toward where you need to be.

Stage 1

Reactive

Security begins at breach detection

Organizations at this stage operate on borrowed time. There is no meaningful visibility into exposure until an attacker has already acted, and every security motion from discovery to remediation is triggered by an incident rather than foresight. Discovery is ad hoc, validation is nonexistent, and prioritization is driven by whichever issue caused the most recent breach. Attacker awareness is effectively zero, and remediation is manual, slow, and always a step behind.

Example: An exposure is only identified after it has already been exploited and appears in an incident response investigation, rather than during any proactive assessment.

Dimensions	Characteristics
Discovery	Ad hoc; assets and vulnerabilities identified mainly through incidents or manual audits
Validation	None; findings treated as risks regardless of actual exploitability
Prioritization	Based on severity scores or whichever issue caused the most recent incident
Attacker Awareness	Generic, based on global trends and point-in-time reports. Minimal direct relevance to organization specifically
Speed of Action	Weeks to months; remediation scramble begins only after an incident occurs
Data Sources & Signals	Internal incident reports, periodic vulnerability scans
Remediation Ownership	Manual, ticket-based
Metrics & Success Measures	Number of incidents resolved; time-to-response
Scope & Coverage	Siloed, asset-by-asset; no consistent attack surface view
Organizational Integration	Security treated as isolated IT function, engaged only post-incident

Stage 2

Compliance-Driven

Ticking the boxes

This stage brings structure, but the structure is built around compliance rather than actual risk. Periodic scans surface vulnerabilities, but validation remains minimal: findings are logged, not tested for real exploitability, and prioritization still leans on static severity scores rather than business context. Attacker awareness is limited to generic threat feeds that aren't tailored to the organization's own attack surface, and remediation happens in backlog-driven bursts tied to audit deadlines.

Example: A vulnerability is flagged during a quarterly scan and sits in a remediation backlog for weeks because it wasn't rated critical, even though it's already being actively discussed by threat actors.

What moves you forward: Replacing periodic assessment with a continuous cycle, and starting to ask which exposures matter most, not just which exist, is the bridge into Stage 3.

Dimensions	Characteristics
Discovery	Periodic vulnerability scans and audits, typically quarterly
Validation	Minimal; vulnerabilities logged but rarely tested for real-world exploitability
Prioritization	Driven by static severity ratings and backlog reduction targets
Attacker Awareness	Limited to generic threat feeds, not tailored to the organization's attack surface
Speed of Action	Weeks (days in certain public sector scenarios for urgent vulnerabilities); remediation cycles tied to regulatory requirements, not exposure urgency
Data Sources & Signals	Scan reports, compliance checklists, generic CVE databases
Remediation Ownership	Manual, backlog-driven; owned by IT/security in isolation
Metrics & Success Measures	Number of vulnerabilities found and patched
Scope & Coverage	Broader than Stage 1 but still point-in-time and incomplete
Organizational Integration	Security engages business only to report top level activity metrics

Stage 3 Continuous

Which exposures matter most?

This is where organizations take their first steps towards processes aligned to Gartner’s Continuous Threat Exposure Management (CTEM) framework, moving from static assessments to an ongoing cycle of scoping, discovery, prioritization, validation, and mobilization. Discovery becomes continuous rather than periodic, prioritization starts factoring in business context, and basic control-effectiveness checks emerge, though systematic exploitability testing isn’t yet in place. Attacker awareness is nascent; exploit intelligence exists but isn’t fully woven into daily decisions.

Example: A critical asset is scoped as high-priority during the CTEM cycle, and its exposures are tracked continuously rather than rediscovered from scratch each quarter.

What moves you forward: Adding rigorous exploitability validation; proving whether an exposure can actually be exploited in your environment, pushes the organization into Stage 4.

Dimensions	Characteristics
Discovery	Continuous identification of vulnerabilities, misconfigurations, exposures and attack paths
Validation	Basic checks on control effectiveness, not yet systematic exploitability testing
Prioritization	Based on business context and severity rather than volume of findings
Attacker Awareness	Emerging; some use of exploit intelligence, not yet integrated into daily decisions
Speed of Action	Days; exposure addressed within an ongoing cycle rather than a fixed calendar
Data Sources & Signals	Internal continuous scanning plus severity scores, business context, and asset criticality
Remediation Ownership	Structured mobilization workflows tied to a repeatable cycle
Metrics & Success Measures	Exposures prioritized and reduced, not just found
Scope & Coverage	Full internal environment plus known external attack surface
Organizational Integration	Security decisions informed by scoping conversations with the business. Ability to report on reduced risk from exposure management activity

Stage 4

Predictive

What is exploitable, and what should we fix first?

The capabilities that underlie Preemptive Exposure Management (PEM) bring real operational muscle to the CTEM framework. Breach and attack simulation and attack path modelling let organizations validate exploitability with more confidence, and prioritization is now driven by simulated attacker behavior rather than theoretical severity. Remediation becomes semi-automated, triggered by validated risk rather than backlog position. The gap that remains: attacker awareness is still simulated, modelling how an attacker could behave rather than tracking what attackers are actually doing.

Example: Simulation confirms that an exposed service is genuinely exploitable and maps a realistic attack path to a critical asset, allowing the team to prioritize it above dozens of higher-severity but less exploitable findings.

What moves you forward: Layering in observable, real-world attacker context; dark web chatter, credential exposure, active targeting — completes the shift into true preemption.

Dimensions	Characteristics
Discovery	Fully continuous and automated across the external attack surface
Validation	Systematic breach and attack simulation (BAS) and attack path modelling to prove or disprove exploitability
Prioritization	Driven by simulated attacker behavior and validated exploitability
Attacker Awareness	Exposures are tied to simulated attacker behavior. Threat intelligence tools are utilized and drive action, though are not fully integrated with exposure data
Speed of Action	Hours to days; validated exposures escalated quickly through automated workflows
Data Sources & Signals	Internal findings, simulation output, exploit intelligence feeds
Remediation Ownership	Fast remediation triggered by validated risk
Metrics & Success Measures	Exploitable exposure reduced
Scope & Coverage	Comprehensive external attack surface, including third-party exposure
Organizational Integration	Exposure reduction reported as a risk metric to leadership

Stage 5

Preemptive

Is there evidence attackers are actively targeting this?

This is the target state, and the one that the PTEM approach is built around. It combines everything from PEM: continuous discovery, exploitability validation, simulated attack paths, with real-world attacker intelligence: exploit development activity, threat actor discussions, credential exposure, and dark web signals. Not threat intelligence as a bolt-on, but attacker reality directly informing prioritization. Exposure windows shrink, remediation is triggered by genuine attacker relevance, and exposure reduction is embedded into business risk conversations rather than treated as a siloed IT metric.

Example: A high priority exposure is prioritized not just because it’s technically exploitable, but has the real-world context attached giving the team confidence to act rapidly before exploitation occurs.

Dimensions	Characteristics
Discovery	Continuous attack surface discovery combined with exposure and security research
Validation	Exploitability confirmed through simulation and proactive research. Instant and continuous verification of closure
Prioritization	Informed by exposure, exploitability, and observable attacker behavior: exploit development, credential exposure, dark web chatter, targeting activity
Attacker Awareness	Real-world and active; asks “is there evidence attackers are actively targeting this?”
Speed of Action	Minutes to hours; exposure windows closed before exploitation occurs
Data Sources & Signals	Comprehensive exposure data. Real time threat intelligence mapped to your organization and its environment, drawn from the clear, deep and dark web, including closed cybercriminal forums.
Remediation Ownership	High confidence findings delivered straight to engineering for remediation
Metrics & Success Measures	Exposure window reduced; exploitable risk eliminated before attackers act
Scope & Coverage	Visibility across your entire external attack surface, including third party software, shadow exposure
Organizational Integration	Exposure reduction embedded into business risk decisions cross-functionally

| Leading the Transformation

Every meaningful shift in cybersecurity has begun with a change in how organizations think. Early threats changed the way we thought about network perimeters. Cloud computing changed the way we thought about infrastructure. Zero Trust challenged long-held assumptions about identity and access. None of these transitions happened overnight, nor were they driven purely by technology. They required organizations to rethink long-established practices, adapt their operating models, and, in many cases, accept that the assumptions which had served them well for years were no longer fit for purpose.

Preemptive cybersecurity is the next fundamental shift.

Throughout this book we have argued that the challenge facing modern security teams is not only a lack of capability, but a mismatch between the way organizations defend themselves and the way attackers actually operate. We have explored how the external attack surface has expanded beyond traditional boundaries, how threat actors increasingly exploit opportunities long before defenders become aware of them, and why frameworks such as PTEM offer a practical way of moving security decisions earlier in the attack lifecycle. None of these ideas stand alone. Together, they describe a different philosophy for approaching cybersecurity, prioritizing visibility, action and prevention over assumptions, reaction and recovery.

Technology undoubtedly plays an important role, but technology alone has never been the thing that transforms organizations. Organizations change because leaders decide to ask different questions, because teams begin working differently, and because success starts to be measured in different ways. The attack surface platform, the threat intelligence capability and the automation that supports them are all enablers of that transformation, but success relies on everything coming together in the right way.

We'll leave you here with an important lesson: becoming preemptive is not about striving for perfection. Every organization represented by the readers of this book will continue to have competing priorities, constrained budgets, technical debt and legacy systems that cannot be modernized overnight. Frontier AI is finding vulnerabilities at an exponential rate, at least until that AI is able to help us build 100% bug-free software. Don't hold your breath. For the foreseeable future, there will not be a point at which an organization can confidently declare that it has eliminated cyber risk altogether. Preemptive cybersecurity has never promised certainty. What it offers instead is a more intelligent way of managing uncertainty.

That means making decisions with better context, understanding your exposure before attackers do, and consistently reducing the opportunities available to those who would seek to exploit your organization. Over time, these decisions compound. An engineering team that considers exposure as part of software delivery, a security team that prioritizes work according to attacker behavior rather than vulnerability volume, and leadership that measures reductions in exposure rather than increases in activity all contribute to an organization that behaves differently.

Collectively these changes create a security posture that becomes progressively harder to attack.


If there is a single theme that runs throughout this book, it is that time has become the defining factor in modern cybersecurity. Attackers have compressed the time between discovery and exploitation. Attack surfaces change in real time. New exposures appear continuously, and organizations can no longer rely on periodic assessments or reactive operating models to keep pace. You shouldn't aim to just 'work faster' within your existing framework. You need to look at what you've got and build a framework that is fundamentally designed around acting earlier.

That is what PTEM represents, providing a way of bringing the right capabilities together around a common objective: understanding where meaningful exposure exists, determining what attackers are most likely to exploit, and reducing that opportunity before it becomes an incident. Sure, it's another acronym in an already crowded space, but considering how threats are rapidly changing, the way of thinking it represents is arguably hard to dismiss.

However long the term PTEM lasts before the next trend arrives is, in many ways, less important than the shift in thinking it represents. The cybersecurity industry has always evolved by challenging assumptions that no longer reflect reality, and we believe the assumption that organizations can continue to defend themselves primarily through reactive security is one of those ideas whose time has passed.

Our hope is that this book has encouraged you to look at your own organization through a different lens. Beyond detecting and responding to attacks, to asking where that attack might begin, whether you would see it coming, and what could be done today to remove the opportunity altogether. Those questions are ultimately what define a preemptive organization, and they are questions that every security leader can begin asking before the next incident forces them to.

New philosophies, techniques and practices in cybersecurity emerge rapidly, and so too do those of the adversaries trying to break them. Adaptability is fundamental. Preemptive cybersecurity is simply the next stage in that evolution. We hope the ideas in this book help you lead that change, so you can spend a little less time responding to yesterday's attacks and a little more time preventing tomorrow's.



PART FOUR

Reference & Learning

The Practitioner's Guide to the Dark Web

For most people, the dark web is synonymous with mystery and criminality. For cybersecurity professionals however, one of the richest sources of insight into attacker behavior available anywhere.

The dark web is where cybercriminals buy and sell access, trade stolen data, discuss vulnerabilities, recruit affiliates, and advertise successful attacks. It is also, with the right visibility, where many of the earliest warning signs of an attack first appear. Long before malware is deployed or compromised credentials are used, there is often evidence of planning, reconnaissance, or preparation taking place in the attacker ecosystem.

Understanding the dark web is therefore not just about curiosity. It is about understanding the environment in which attackers operate, so defenders can identify threats earlier and act before they become incidents.

What is the Dark Web?

The internet is commonly divided into three layers. The surface web is the small portion indexed by search engines like Google. The deep web consists of content that is not publicly indexed, such as corporate applications, databases, and password-protected websites. The dark web is a subset of the deep web that requires specialist software, most commonly the Tor network, to access.

The dark web was originally designed to provide anonymity and protect privacy. Today, while it still has legitimate uses for journalists, activists, and whistleblowers, it has also become home to a thriving cybercriminal economy. Hidden services hosted on Tor provide forums, marketplaces, messaging platforms, and leak sites where threat actors can operate with a greater degree of anonymity than on the public internet.

It is important to remember that the dark web is not the entirety of the attacker ecosystem. Modern threat intelligence also draws on encrypted messaging platforms such as Telegram, closed criminal forums, paste sites, code repositories, and other parts of the deep and clear web. However, the dark web remains one of the most valuable places to observe cybercriminal activity as it unfolds.

Why the Dark Web Matters

From a preemptive security perspective, the value of the dark web is simple: it provides visibility into the attack lifecycle before an attack reaches your organization.

Attackers rarely move directly from identifying a target to deploying ransomware or stealing data. First they gather intelligence, purchase stolen credentials, discuss vulnerabilities, advertise compromised networks, recruit partners, and exchange tools. Much of this activity takes place on dark web forums, marketplaces, and ransomware infrastructure.

These environments generate valuable signals, including:

- Stolen employee credentials appearing for sale.
- Initial Access Brokers advertising access to an organization's network.
- Threat actors discussing newly discovered vulnerabilities or exploits.
- Ransomware groups publishing victim data or recruiting affiliates.
- Mentions of specific organizations, sectors, or technologies as potential targets.

If attackers are actively discussing a vulnerability affecting technology you use, or advertising access to your organization, that exposure moves immediately to the top of the remediation queue.

This is the essence of the threat intelligence layer within PTM. Rather than reacting after malicious activity reaches your network, defenders monitor the environments where attacks are planned and prepared, shifting detection further left in the attack lifecycle and buying valuable time to act.

The Dark Web Ecosystem

The dark web consists of a diverse ecosystem of criminal platforms, each serving a different purpose. Together, they provide a valuable view into how attackers communicate, collaborate, monetize stolen data, and plan attacks. The categories below highlight some of the most important parts of that ecosystem.

Hacking Forums

Dark web forums are very similar to their clear web counterparts in almost every way, except for the content they discuss. Indeed, some began life on the clear web before migrating to the dark web to evade surveillance or censorship, though others continue to maintain a clear web presence.

Forums have much greater longevity than dark web marketplaces. Some of the most notorious Russian-language forums (for example, XSS and Exploit) have been around for decades.

BF

BreachForums

ACTIVE SINCE: March 2022

DARK WEB OR CLEAR WEB: Both

PREDOMINANT LANGUAGES: English

KNOWN ALIASES OF ADMINS OR STAFF: Baphomet, ShinyHunters, Manitoba, Dedale

BreachForums launched in March 2022 as the successor to RaidForums, operating on both the clear and dark web to trade leaked databases and cybercrime tools. It first shut down in March 2023 after the arrest of administrator “pompompurin,” but quickly re-emerged in June 2023 under veterans like “Baphomet” and the hacking collective “ShinyHunters”.

Despite an FBI seizure in May 2024, the forum returned just two weeks later, highlighted by a massive Ticketmaster database leak. Management underwent rocky changes in July 2024 when ShinyHunters stepped down for a poorly received new owner, Anastasia. By April 2025, the forum mysteriously shut down again.

After several failed revival attempts, the site briefly returned in December 2025 by restoring an old database backup. However, this latest iteration was short-lived; by June 2026, it shuttered its operations entirely, with its administrators confessing they were “just another clone” impersonating the original ShinyHunters.

A new version of BreachForums surfaced in August 2026, claiming not to be associated with previous iterations.



Cracked

ACTIVE SINCE: April 2018

DARK WEB OR CLEAR WEB: Clear Web

PREDOMINANT LANGUAGES: English

KNOWN ALIASES OF ADMINS OR STAFF: KSZ, Barry, Darkness, Liars, Ping

Cracked is a clear web hacking forum that takes its name from the act of “cracking”, slang for breaking into accounts or software (usually with the intention of circumventing payment).

Users trade in tools, configs, tutorials and other resources to achieve this end, such as leaked credential combo lists and SOCKS proxies, as well as discussing how to monetize their illicit activities.

Cracked was seized in January 2025 in Operation Talent, a German-led law enforcement operation that also disrupted Nulled forum. Cracked has since returned virtually unchanged from its pre-seizure form.



DarkForums

ACTIVE SINCE: November 2022

DARK WEB OR CLEAR WEB: Both

PREDOMINANT LANGUAGES: English

KNOWN ALIASES OF ADMINS OR STAFF: Lokii, Tanaka, Solid Snake, Zexxy, Knox

A spiritual successor to the RaidForums/BreachForums lineage, DarkForums has grown to become the top destination for selling and sharing stolen databases, at least in the English-speaking cybercrime sphere. Originally launched in 2022 as a spin-off of Indian hacking forum D4RK4RMY, DarkForums gained more and more users as the continuous seizures of BreachForums iterations took their toll on the brand's perceived reliability.



Dread

ACTIVE SINCE: February 2018

DARK WEB OR CLEAR WEB: Dark Web

PREDOMINANT LANGUAGES: English

KNOWN ALIASES OF ADMINS OR STAFF: HugBunter, Paris, Syntax, Shakybeats, Solar

Dread is a dark web forum that was born out of Reddit's clampdown on discussions around dark web markets and scamming techniques.

Almost since its inception it has been plagued by denial-of-service (DDoS) attacks that have at times left it virtually unusable. According to chief administrator HugBunter, the downtime is the result of a persistent adversary targeting the hidden service with the objective of extorting dark web markets that use Dread to communicate with their users. This downtime has led many users to migrate to Dread's I2P mirror to access the dark web forum.

Nevertheless, it remains a popular hub for dark web netizens involved in the market scene, as well as those interested in broader cybercrime. Conversations on Dread are mainly focused on market and vendor reviews (with many dark web users relying on Dread as a source of information) but it also has popular subs on fraud techniques such as carding, hacking, and cybercrime job recruitment.



Exploit

ACTIVE SINCE: February 2005

DARK WEB OR CLEAR WEB: Both

PREDOMINANT LANGUAGES: Russian

KNOWN ALIASES OF ADMINS OR STAFF: Admin, Support, Garant, Adv, BigBear, L.Luciano, Mr.Burns, Pixel, JohnRipper, Oxygen, Quake3, Weaver

Exploit is an extremely long-running Russian cybercrime forum that has been active since at least 2005.

Exploit and other Russian forums tend to view themselves as more professional than other dark web communities, often shunning non-Russian speakers and those perceived as unskilled or inexperienced.

As such, the site acts as something of a network for career cybercriminals to connect with potential collaborators on illegal business ventures, be it hacking, scamming, or working on Ransomware-as-a-Service (RaaS) schemes.

At times the ransomware aspect has been tempered – for example, during the unwanted attention Exploit received in the wake of the 2021 Colonial Pipeline attack.

In this vein, we regularly witness threat actors auctioning initial access to organizations, usually through VPN or other remote access software. The posts typically have a “start” price to kick off the auction, a “step” price that indicates the increments of bidding, and a “blitz” price if a bidder wants to buy the access outright.



RuTor

ACTIVE SINCE: October 2014

DARK WEB OR CLEAR WEB: Both


PREDOMINANT LANGUAGES: Russian

KNOWN ALIASES OF ADMINS OR STAFF: Senior Berlin, Screaming Eagle, Darkless, Viktor Palych

RuTor has been a prominent forum in the Russian dark web scene since 2015, with threads that span multiple pages – and years – on topics ranging from drugs, fraud, computer programming and current affairs.

It is known for hosting region/city-specific “smoking rooms,” which acted as assembly points for dispossessed vendors and buyers following the takedown of drug market giant Hydra in April 2022.

Hydra’s demise, coupled with the Russian invasion of Ukraine, triggered a civil war of sorts in the Russian cybercrime community. RuTor suffered a hack at the hands of pseudo-hacktivist collective Killnet, before launching its own attack against rival forum WayAway.



~/XSS.is

XSS

ACTIVE SINCE: November 2004

DARK WEB OR CLEAR WEB: Both

PREDOMINANT LANGUAGES: Russian

KNOWN ALIASES OF ADMINS OR STAFF: Admin, Guron_18, Haunt, IIIIXX, Kerberos, Marcus52, Pernatly, Quake3, R_as, Weaver

Originally known as DaMaGeLaB, XSS is one of the longest-running dark web forums.


The site rebranded from DaMaGeLaB to XSS around 2018, potentially due to the arrest of one of its administrators a year prior for their involvement in operating the Andromeda botnet. Its new name is a reference to the well known cross-site scripting web app vulnerability.

As with Exploit, XSS is very business-oriented, with sections on hacking, corporate access, database leaks, and even competitive intelligence.

XSS had previously acted as a recruitment and PR tool for Ransomware-as-a-Service (RaaS) schemes, although this content has been “banned” at certain times, presumably so the forum doesn’t attract too much unwanted attention from law enforcement.

In 2023 the forum appeared to trial an “XSSBot”, a forum chatbot that we suspect used ChatGPT to power its responses. Forum users asked the XSSBot for information about different malware strains, for tips on how to obfuscate code, and to write a rap – among other things.

The suspected administrator of Xss was arrested in Ukraine in July 2025, in a French-led law enforcement operation. This led to several splinter forums being created, though Xss itself does continue to operate.



Ransomware Leak Sites

Ransomware groups use the dark web for multiple stages of their cyberattacks.

Many use some of the dark web forums we have already discussed to interact with the wider cybercriminal community. In particular, to buy initial access to organizations that they can use to gain a foothold in their victims' networks. Or – if they are a Ransomware-as-a-Service (RaaS) operation – they may use forums to recruit affiliates.

However, the key presence of ransomware groups on the dark web are leak sites – which are used by dozens of threat actors (we currently track more than 100 and this number is always increasing).

Ransomware leak sites are effectively publicity sites where ransomware groups share the details of their latest victims. However, they also play an important role in how these groups orchestrate and monetize their attacks. These sites provide the ransomware operators with a platform to accept payments from the victims, a space to shame them and apply pressure, and somewhere to leak their data if they don't pay.



AKIRA

ACTIVE SINCE: March 2023

TOP TARGETED GEOGRAPHIES: US, UK, Canada

Akira is a prolific ransomware group operating under the RaaS business model. The group was first identified in March 2023 and claimed at least 1200 victims (as of the end of 2025) via their extortion site.

The ransomware variant employed by the group has evolved over time, with first strains being developed using the C++ programming language and giving files the .akira extension when encryption was completed. The group has started using the Megazord ransomware variant as well, which uses the .powerranges extension.

Like many other ransomware groups, Akira's affiliates obtain initial access from a multitude of sources, such as stealer logs or Initial Access Brokers. Access is also obtained by exploiting technical vulnerabilities found in common commercial solutions such as Cisco or SonicWall VPN software and remote desktop protocol (RDP) apps. Social engineering such as phishing and spearphishing are also part of their tactics.

Once an initial foothold is established, the group often employs an attack technique known as Kerberoasting to dump the Local Security Authority Subsystem Service (LSASS) and obtain additional credentials. Mimikatz and LaZagne have also been used to obtain additional credentials. Exfiltration is done by leveraging legitimate tools including FileZilla, Rclone, WinSCP and WinRAR.

Although Akira is likely a self-established group with no major links to other ransomware groups, past or present, there are some indications that certain elements of the malware itself were inspired by Conti. Moreover, some cryptocurrency payments can be traced to former members of the Conti affiliate program.

The ransomware strain is sophisticated and uses a hybrid encryption method, however, a decryptor was created and shared in the community on two different occasions, one in 2023 and one in early 2025. The latter was particularly interesting as it uses GPUs to bruteforce to decrypt keys. However, it only worked on Linux-based systems and required a significant upfront investment to purchase numerous GPUs. It was also a lengthy process to obtain a decryption key, it took 16 high end GPUs 10 hours to decrypt one single key.

CLOP^ _ -

CLOP

ACTIVE SINCE: February 2019

TOP TARGETED GEOGRAPHIES: US, UK, Canada

CLOp ransomware is known to be used by the cybercriminal enterprise tracked as TA505 and FIN11.

CLOp is notable for its approach of using vulnerabilities in supply chain software to target multiple organizations, announcing them in a batch at a later date. This was a tactic it used to great effect into 2023, with two “mass-hacks” making CLOp the third most prolific ransomware group of the year by number of listed victims (after LockBit and BlackCat).

In March of that year, CLOp exploited the vulnerability CVE-2023-0669 in Fortra’s GoAnywhere MFT secure file transfer tool to target more than 130 organizations, listing them in quick succession. Then in June, CLOp repeated this approach in one of the biggest and most notable cyberattacks of the year, exploiting a zero day vulnerability (CVE-2023-34362) in the Progress Software file transfer software tool, MOVEit.

The group had so many victims from the MOVEit breach that it had to explore new ways of leaking data, including using torrents. While hundreds of companies were listed on its leak site, it was reported that there were in fact more than 1,000 organizations impacted by the MOVIEit attacks. Noteworthy victims included the BBC, British Airways, Emsisoft, U.S. government services contracting company Maximus, and the French government's unemployment agency, Pôle emploi.

In the aftermath, the U.S. State Department offered a \$10 million bounty for information on CLOp and the group's activity has fluctuated since then.



DragonForce

ACTIVE SINCE: November 2023

TOP TARGETED GEOGRAPHIES: US, UK, Germany

DragonForce emerged in late 2023, and since then has repeatedly proved itself as a force to be reckoned with.

Originally a hacktivist operation, DragonForce pivoted to a RaaS model and has since innovated on the concept, with its “ransomware cartel” program in which threat actors can use its malware under their own separate branding. This offer was put to use with destructive effect in the attacks on British retailers Marks & Spencer, Co-op and Harrods in Spring 2025, which were attributed to the Scattered Spider threat group using DragonForce ransomware. DragonForce has even used adversarial tactics against its competitors, seemingly being implicated in the disintegration of at least two other RaaS operations (RansomHub and BlackLock).

In September 2025, on a closed Russian hacking forum, DragonForce announced a coalition with other RaaS operators, namely Qilin and LockBit. The proposed aim of the coalition was to pool resources and work together to increase overall income. Since this announcement, as evidenced by our data, Qilin's victim count and prominence increased, potentially demonstrating another example of the success of the 'Super Group' collaboration model.



EVEREST GROUP

Everest

ACTIVE SINCE: December 2020

TOP TARGETED GEOGRAPHIES: US, Canada, Italy

Everest has been around since at least 2020, making it one of the oldest ransomware operations still active after LockBit and ClOp.

Initially using the now-popular technique of double extortion – encrypting a victim’s data in addition to stealing and threatening to publish it on its dark web leaks site – Everest has claimed in recent years to eschew ransomware and engage in data extortion-only attacks. The gang has also been observed moonlighting as an initial access broker, providing unauthorized corporate network access to other threat actors for a fee.

INC Ransom

INC Ransom

ACTIVE SINCE: June 2023

TOP TARGETED GEOGRAPHIES: US, Canada, UK

IncRansom is a financially motivated group that emerged in 2023. The group have operated a RaaS model and so victims are spread across sectors/industries and geographically. As with many ransomware groups, double extortion, the encrypting, downloading and applying pressure on the victim is the general tactic utilized. Again, similar to many RaaS operators, a mixture of initial access vectors are exploited, ranging from using purchased credentials, conducting phishing campaigns or exploiting known vulnerabilities.



Qilin

ACTIVE SINCE: Mid 2022

KNOWN ALIASES: Haise

TOP TARGETED GEOGRAPHIES: US, Canada, UK

Qilin is a prolific ransomware group using the RaaS business model. The group became active in mid 2022, initially using the name “Agenda”. At the time of writing this report, the group has claimed over 1,300 victims via its extortion site.

The initial ransomware variant was developed using the Go programming language, however, this was rewritten using Rust and later C. Variants target systems using the Windows, Linux and ESXi operating systems. The group is known to heavily rely on social engineering techniques to gain initial access.

A member of the Ramp cybercrime forum using the ‘Haise’ handle is a representative of the ransomware group and often advertises the affiliate program with the intention of recruiting new members.

The forum user claimed that the ransomware is capable of delivering four types of encryption, encrypting a file in full or in part. Successful candidates for the RaaS scheme are granted access to a panel featuring the build configurator, dialogue support, and spam calling and SMS services.

Qilin made several updates over the course of 2025, and added an option to conduct distributed denial-of-service (DDoS) attacks against victims to pressure them into paying the ransom. This is known as triple extortion. Additionally, affiliates are offered legal advice on how to negotiate with victims based on their jurisdiction and type of compromised data, tailoring, and streamlining their negotiation process to each victim.

In the latest update, the Qilin spokesperson stated that a call center available in seven languages would be soon opened to contact victims and their customers. All of this suggests an acceleration of activity and – indeed – compared to the other most active groups, Qilin has generally posted more victims month-on-month.

Dark Web Marketplaces

Dark web marketplaces are anonymous online trading platforms that function much like legitimate e-commerce sites, complete with vendor ratings, escrow services, customer reviews, and dispute resolution mechanisms. While many facilitate the sale of illegal drugs, counterfeit goods, forged documents, and other forms of organized crime, they have also become an important part of the cybercriminal economy.

Escrow Marketplaces

Escrow marketplaces allow anyone with enough cash to pay the “vendor bond” to sign up and start selling. These are like the dark web equivalents of Amazon or eBay, except with less reputable products and payments made using cryptocurrencies.

Escrow marketplaces are a particularly dynamic part of the dark web. Many of the goods sold on them (drugs, stolen credit cards, fake IDs, counterfeit goods, malware, etc.) are explicitly illegal and the operators therefore find themselves under huge pressure from law enforcement.

Even those that avoid being seized are aware that their days are numbered, which means they are liable to close up shop at any moment to protect themselves (often stealing the held funds of their vendors or buyers in the process).

BLACK OPS

BlackOps

ACTIVE SINCE: September 2024

DARK WEB OR CLEAR WEB: Dark Web

APPROX. NUMBER OF LISTINGS SINCE LAUNCH: 18.2K

APPROX. NUMBER OF VENDORS SINCE LAUNCH: 893

PREDOMINANT LANGUAGE: English

BlackOps is a fairly new dark web market, beginning activity in September 2024. It quickly gained popularity, perhaps owing to unique features such as a variable vendor fee based on number of sales and the option to add market users as employees. Its product offerings are fairly standard, though one quirk is that the number of digital products for sale exceeds its drug-related listings.

DARK MATTER

Darkmatter

ACTIVE SINCE: September 2022

DARK WEB OR CLEAR WEB: Dark Web

APPROX. NUMBER OF LISTINGS SINCE LAUNCH: 37.6K

APPROX. NUMBER OF VENDORS SINCE LAUNCH: 1.6K

PREDOMINANT LANGUAGE: English

KNOWN ALIASES OF ADMINS OR STAFF: quasar1, BlackMask, zamman

While it's considerably smaller in size than other markets featured, DarkMatter has been consistently building a user base while avoiding the typical scandals that often plague its larger peers.

DarkMatter was launched in September 2022, and differentiates itself by only allowing Monero (rather than the older and less privacy-focused Bitcoin) transactions. It also offers a walletless payment system, meaning customer deposits aren't stored on the site, in an attempt to assuage fears of exit scams. Predominantly hosting drug vendors, DarkMatter also offers a number of fraud, counterfeit, and malware products.

Drughub

DrugHub

ACTIVE SINCE: August 2023

APPROX. NUMBER OF LISTINGS SINCE LAUNCH: 113.1K

APPROX. NUMBER OF VENDORS SINCE LAUNCH: 2.5K

PREDOMINANT LANGUAGE: English

KNOWN ALIASES OF ADMINS OR STAFF: DrugHub, DrugHubStaff, DrugHubStaff2

DrugHub's creators claim to be former staff members of WhiteHouseMarket, an extremely popular dark web market which closed its doors (without exit-scamming) in October 2022. Trading on this good reputation appears to be working, as DrugHub has become a frontrunner in both number of listings and number of active vendors in recent months. Key features include Monero-only transactions, mandated PGP login, and a multi-vendor shopping cart system.

DrugHub recently absorbed the userbase of another platform, SuperMarket, after one of SuperMarket's admins was alleged to have absconded with its customers' on-site cryptocurrency.



Torzon

ACTIVE SINCE: September 2022

APPROX. NUMBER OF LISTINGS SINCE LAUNCH: 80.9k

APPROX. NUMBER OF VENDORS SINCE LAUNCH: 1.3k

PREDOMINANT LANGUAGE: English

KNOWN ALIASES OF ADMINS OR STAFF: TorZon, Kakarot, tOnka, Nintendo

Torzon launched in September 2022. As of July 2025 it has resided on darknet forum Dread's superlist, an almanac of dark web markets deemed trustworthy by the community of users. Its primary offerings are standard: drugs, fraud, counterfeits and digital products.

Autoshop Marketplaces

Autoshops are a particular type of dark web marketplace that specialize in the sale of digital products. They differ from escrow marketplaces in two main ways. Firstly, the transaction is automated, meaning that there is little to no contact with the seller. Secondly, there are usually far fewer vendors active on the site, suggesting a more exclusive selection process than the “open application” ethos of escrow marketplaces.

Autoshops specialize in the sale of information, such as financial data, login credentials, remote access, and cookies, which means they typically have a high turnover of listings – sometimes into the millions. Their administrators also seem to prefer keeping a lower profile than some of their escrow market counterparts.

Autoshops are also more likely to operate on the clear web but often have some form of entry barrier, such as requiring a minimum account balance or an invite from a known user. In the case of invite-only shops, staff will sometimes distribute invites on hacking forums or existing customers generate and sell them on forums or messaging apps such as Telegram.

b1ack's Stash

BlackStash

ACTIVE SINCE: January 2024

DARK WEB OR CLEAR WEB: Both

BlackStash is a carding autoshop that opened in early 2024, but its admin was actively sharing large datasets on cybercrime forums 6 months earlier. Despite successful monetization, BlackStash still promotes their store by regularly releasing batches of free data on predominantly Russian-speaking forums, including a recent drop of over 4 million stolen credit card details.



Bahira

ACTIVE SINCE: March 2022

DARK WEB OR CLEAR WEB: Both

Fraudsters began recommending Bahira Finance on dark web forums in early 2022. The name, which is shared by a Christian monk in Islamic theology, was likely selected due to its meaning of “tested” or “approved” in the Syriac language. Fittingly, Bahira offers numerous validity and balance checkers for the card details it sells, which is important to its customers who often bemoan the common experience of purchasing “dead” cards.

Bahira has generated buzz in both Russian and English-speaking dark web spaces and boasts the sale of both card details and dumps (the information encoded onto a card’s magnetic stripe and used to make physical clones). Unlike some autosshops, users can apply for vendor accounts on the site, rather than being part of a closed group or selling card databases behind the scenes.



BlackPass

ACTIVE SINCE: May 2017

DARK WEB OR CLEAR WEB: Both

BlackPass specializes in stolen login details needed to hijack e-commerce accounts – which often have associated payment data – rather than the card details directly. While all listings include at least the credentials required to log into accounts for sites like Amazon, Venmo, and Instagram, some accounts have PII associated with them such as the victim's full name, country, ZIP code, bank and card details, and phone number.

BlackPass closely resembles Slilpp, another popular credential autoshop disrupted in a June 2021 multinational law enforcement operation.



BriansClub

ACTIVE SINCE: 2014

DARK WEB OR CLEAR WEB: Both

With earliest estimates placing its inception at 2014, BriansClub is the oldest autosshop on this list. It is well known in the cybercriminal underground, not least for its hijacking of security journalist Brian Krebs's likeness.

BriansClub sells a range of fraud products, including CVVs, fullz (card details packaged with additional cardholder information such as date of birth and social security number), and dumps.

BriansClub also has a wholesale section, selling full databases of stolen card details for bulk buyers.



PatrickStash

ACTIVE SINCE: March 2022

DARK WEB OR CLEAR WEB: Both

PatrickStash began being mentioned on cybercrime forums as a viable option for carders in March 2022.

The site's forum representative is very active in advertising the shop, touting automatic refunds, live statistics, and sellers in a range of countries as their USPs. On this note, the PatrickStash forum account regularly puts out recruitment calls for new vendors to sell to them, garnering a generally positive response.

PatrickStash has two categories, Cards and Cards NoVBV (short for Verified By Visa, now called 3D-Secure, an additional security layer meant to protect against online payment fraud). Despite its overtly Irish theme, the site is more likely to originate from Russia, or at least targets its advertising towards that audience.



RussianMarket

ACTIVE SINCE: January 2019

DARK WEB OR CLEAR WEB: Both

RussianMarket specializes in the sale of “logs” – records harvested by information-stealing malware – as well as CVVs, dumps and RDP (remote desktop protocol) access.

Unsurprisingly, it is suspected to be of Russian origin. Indeed, many of the examples of autoshop marketplaces we have listed are assumed to be of Russian origin, based on the language preferences of the sites and the forums used to promote their wares.

Autoshops are also more likely to operate on the clear web but often have some form of entry barrier, such as requiring a minimum account balance or an invite from a known user. In the case of invite-only shops, staff will sometimes distribute invites on hacking forums or customers sell them as a third-party broker on forums or messaging apps such as Telegram.

| PTEM Buyers Guide

Choosing Technology That Enables Preemptive Cybersecurity

Throughout this book, we've explored why cybersecurity must evolve from reacting to attacks towards reducing exposure before attackers can exploit it. We've examined the changing threat landscape, the limitations of traditional vulnerability management, the importance of understanding your true attack surface, and the role of attacker intelligence in making better security decisions. You've seen how, with the right capabilities, reacting to incidents is not inevitable.

The next question is of course: What technology do I need to make this possible?

Unfortunately, the market doesn't make answering that question easy.

Attack Surface Management vendors increasingly describe themselves as Preemptive platforms. Threat Intelligence providers claim to deliver exposure management. Vulnerability management vendors promise attack surface visibility, while BAS, EASM, CAASM, and CTEM vendors all compete for similar budgets using increasingly similar language.

Many of these products solve genuine problems, but very few deliver every capability required to support a truly preemptive security program.

This guide is designed to help you evaluate technologies through the lens of your preemptive security goals rather than a single product category.

Instead of asking:

“What’s the best ASM platform?” or “Which Threat Intelligence vendor should we buy?”, keep in mind the ultimate question:

Will this technology help our organization become more preemptive?

Start with Outcomes

One of the biggest mistakes organizations make when evaluating cybersecurity technology is beginning with product categories instead of security outcomes.

Instead of starting with product, start with these questions:

1. Can we continuously discover new exposures?
2. Can we determine which exposures are genuinely exploitable?
3. Can we understand whether attackers are actively targeting them?
4. Can we prioritize remediation based on real-world evidence of imminent threat?
5. Can we reduce exposure before attackers exploit it?

If a technology cannot help answer most of these questions, it is unlikely to support a genuinely preemptive security program.

You're Buying a Capability

PEM, as coined by Gartner, or PTEM as we have outlined, is not another individual product category, but an operating model.

- Attack Surface Management
- Threat Intelligence
- Exposure Validation
- Attack Path Analysis
- Breach and Attack Simulation

These are all valuable capabilities, but none of them, individually, constitute PTEM.

PTEM is the operational integration of continuous discovery, exploitability validation, attacker awareness, intelligent prioritization and rapid remediation.

When evaluating vendors, ask how well they enable each capability required to operate preemptively.

The Six Capabilities Every PTEM Platform Should Provide



Continuous Attack Surface Visibility

LOOK FOR

- Internet-facing infrastructure
- Cloud assets
- APIs
- DNS infrastructure
- SaaS applications
- Third-party exposure
- Shadow IT
- Shadow Exposure
- Forgotten assets
- Newly deployed services

ASK THE VENDOR

- How frequently are new assets discovered?
- What constitutes an “asset”?
- Can discovery occur continuously?
- How are cloud-native environments handled?
- Can temporary or ephemeral assets be identified?

ASK THEM TO DEMONSTRATE

- Discovery of a previously unknown asset
- Continuous monitoring
- Identification of cloud assets
- Discovery of shadow exposure

RED FLAGS

- “We perform daily scans.”
- “We maintain an asset inventory.”

Everything begins with visibility. If you cannot continuously discover your attack surface, you cannot continuously reduce it.

An effective PTEM platform should provide continuous visibility across your external attack surface, rather than relying on periodic scans or manually maintained inventories.



Exposure Validation

Not every vulnerability deserves immediate attention. PTEM is about identifying exposures attackers can actually exploit. The platform should distinguish theoretical risk from validated risk.

LOOK FOR

- Proof-of-Concept validation
- Attack path analysis
- Exploitability testing
- Misconfiguration validation
- Control effectiveness testing
- High-confidence findings

ASK THE VENDOR

- How do you determine exploitability?
- Do you generate PoCs?
- How are false positives reduced?
- What evidence accompanies every finding?
- Is exploitability continuously revalidated?

ASK THEM TO DEMONSTRATE

- A validated exploit
- Evidence behind an alert
- False-positive reduction
- Attack-path validation

RED FLAGS

- "We alert on every CVE."
- "We prioritize by severity."
- "We rely on CVSS."



Attacker Activity

This is what distinguishes PTEM from conventional exposure management. Knowing an exposure exists is useful; knowing attackers are actively targeting it changes everything.

LOOK FOR

- Closed cybercriminal forums
- Dark web marketplaces
- Credential leaks and infostealer logs
- Initial Access Brokers
- Ransomware leak sites
- Telegram
- Exploit communities and hacking forums
- Clear web security research and OSINT

ASK THE VENDOR

- Where does your intelligence originate?
- How much is proprietary?
- How frequently is it collected?
- How quickly is it made available?
- Is the intelligence delivered relevant to my organization?
- Is the intelligence clear and actionable?
- Does intelligence directly influence prioritization?

ASK THEM TO DEMONSTRATE

- How real-time attacker activity is pulled into the platform
- Relevant credential exposure
- Threat actor profiling
- How intelligence changes remediation priority

RED FLAGS

- “We integrate commercial threat feeds.”
- “We provide static reports.”
- “We publish monthly intelligence.”



Intelligent Prioritization

Security teams rarely suffer from too few findings, and instead suffer from too little confidence about which findings matter most.

PTM combines multiple signals to prioritize work according to genuine attacker risk.

LOOK FOR

- Business criticality
- Exploitability
- Attack paths
- Exposure age
- Active exploitation
- Threat actor targeting and activity
- Credential exposure
- Zero-day research

ASK THE VENDOR

- What determines priority?
- Is CVSS only one input?
- How much weight comes from attacker activity?
- Can prioritization be customized?
- How is business context incorporated?
- When does prioritization happen?
- Do I get alerted to everything or just what matters?

ASK THEM TO DEMONSTRATE

- How attacker activity changes risk
- Custom prioritization

RED FLAGS

- “We sort by CVSS.”
- “We use AI.”... ask what inputs the AI actually uses.



Proprietary Research

One capability increasingly separates market leaders from everyone else. Research.

Preemptive security depends upon knowing things before everyone else does.


Ask whether the vendor actively contributes new knowledge to the industry, or simply republishes everyone else's.

ASK THE VENDOR

- Do you employ offensive security researchers?
- Have you discovered new vulnerabilities?
- Do you publish original research?
- Do you develop Proofs of Concept?
- Do you contribute CVEs?
- What percentage of your intelligence is proprietary?

Why it matters:

The earlier a vendor discovers emerging vulnerabilities and attacker behavior, the earlier your organization can respond.





Operationalization & Remediation

Discovery without action has little value. PTEM should measurably reduce exposure, not simply identify it.

LOOK FOR

- Workflow automation
- Engineering integrations
- Ticketing
- SIEM
- SOAR
- APIs
- Remediation tracking
- Continuous verification

ASK THE VENDOR

- How quickly can findings be remediated?
- Is remediation verified automatically?
- What integrations exist?
- How much manual effort is required?

ASK THEM TO DEMONSTRATE

- Integration with your existing security workflows
- Remediation verification
- Automatic rescanning

Common Preemptive Myths

As preemptive cybersecurity solutions become more widely adopted, expect many vendors to adopt the terminology.

Always test marketing claims against operational capability.

“We support PTEM because we integrate Threat Intelligence.”

Integrating generic threat feeds does not provide the actionable context to close the exposure window before today’s attackers strike

“We prioritize using CVSS.”

If you’re having to retroactively prioritize an avalanche of theoretical risk, your tool is causing you more problems. And remember, severity is not real risk.

“We scan every day.”

Exploitation timeframes are now measured in hours. Don’t give attackers a 23 hour head start.

“We use AI.”

Ask what signals feed the model.

“We support CTEM.”

Which stages? How?

“We have threat feeds.”

Do they actually change remediation priorities?

Questions Every Buyer Should Have Answered

By the end of every vendor demonstration, you should know:

Can they discover my entire attack surface in real-time, not just once a day?

Can they prove exploitability?

Can they show me what attackers are actually doing?

Can they prioritize using attacker context?

Can they integrate into my remediation workflows?

Can they continuously verify remediation?

Can they demonstrate these capabilities rather than simply describe them?

Can they give me a measurable head start over attackers?


If the answer to any of these questions is unclear, keep asking.

Remember

Technology alone will not make an organization preemptive. Preemptive Threat Exposure Management is an operating model that combines continuous visibility, exploitability validation, attacker awareness, intelligent prioritization and rapid remediation into a single, ongoing process.

The right technology can dramatically accelerate that journey, but success ultimately depends on how those capabilities are integrated into your people, processes and decision-making.

Don't buy a platform simply because it claims to be preemptive. Buy based on the capabilities that help your organization close the exposure window before attackers act, and deliver the ultimate outcomes you need.



| Glossary of Key Terms

Threat Intelligence

Account Takeover (ATO) – An attack technique where a threat actor takes over a legitimate user account using stolen passwords and usernames, which have often been purchased from the dark web.

Active Directory – An internal directory service (think digital phone book) that allows administrators to manage user permissions and control user access to network resources. Developed by Microsoft for Windows domains.

Advanced Persistent Threat Group (APT)
– Groups that carry out attacks on a nation's cybersecurity or economic assets through acts of cyberespionage or sabotage.

Autoshop – A type of dark web marketplace that specializes in the sale of credit card, debit card or bank account information, as well as the credentials, cookies and remote access needed to take over online accounts. The term “autoshop” refers to the transaction process being automated. Autoshops sell digital products, which means that the purchased item can be delivered instantly to the buyer with little to no input required from the vendor.

Backdoor – Refers to a concealed vulnerability or exploit in a system that can be used to bypass security features and gain access to restricted data. This can be added on purpose by developers for later access or by accident, for example by leaving a testing password in the release version that was intended to be removed.

Botnet – A collection of infected computers that form a network, which is remotely controlled by a threat actor. Botnets can be used to perform coordinated functions, combining their computing power to conduct DDoS attacks or sending spam emails. Botnets attacks originate from many different IP addresses, making them difficult to trace.

Browser fingerprints – Browser Fingerprints are harvested from computers that have been infected with information stealer malware. They often include data such as credentials, payment information, cookies, IP addresses and user-agent strings, which can be used to bypass anti-fraud solutions by making the browser session appear identical to the victim's.

Brute Force Attack – A technique where the attacker uses computing power and time to gain access to a network. These attacks often follow a loose set of rules. For example, in the case of brute-forcing a password, an attacker might start at 0 and progress incrementally trying each possible combination before moving onto the next option.

Canary – A PGP-signed message, usually posted by dark web market or forum administrators to signify that they are still in control of the site (not to be confused with a warrant canary, which is a method used to alert people that a gag order has been served on a site or provider).

Carding – The process of stealing personal and banking card details. These are either sold to the highest bidder or used for identity fraud, to extract as many funds as possible from an individual's accounts.

Clear Web – The clear (or surface) web comprises all the publicly available websites that are indexed in search engines, where anyone can find and access them.

Combos – Combinations of credentials, usually names, emails, and passwords, that all relate to a single individual.

Command and Control (C2) Server – Cybercriminals use command-and-control servers to remotely control compromised devices. This server is used to move laterally through systems and exfiltrate stolen data from an affected network.

Credential Stuffing – Using credentials from previous cyberattacks or leaks to try and gain access to a different account linked to the same user.

Critical National Infrastructure (CNI) – Infrastructure that is vital to the functioning of a society or a state, including water supply, transport, networks, and energy companies.

Cross-Site-Scripting (XSS) – Cross-Site-Scripting (XSS) is a cyberattack technique where malicious code is “injected” into trusted, legitimate websites.

Crypter – A type of software that can encrypt, obfuscate, and manipulate malware, making it more difficult for security programs to detect it. Threat actors use it to create malware that can bypass security programs as it appears to be a harmless program until it is installed on the victim's system.

Cryptocurrency – A digital currency used as an alternative form of payment. Usually built on the blockchain and created using encryption algorithms, with an emphasis on anonymity.

Cryptocurrency Exchange – A business that allows users to exchange different cryptocurrencies or change their cryptocurrencies for traditional (fiat) currency such as USD.

Cyber Kill Chain – Originally developed by Lockheed Martin, the Cyber Kill Chain is a framework to explain the series of steps a cybercriminal must complete in order to execute their attack.

Dark Web – The dark web is a subset of the deep web, which is deliberately obfuscated and often requires specialized software, such as Tor, to access. It accounts for approximately five percent of the internet and, while it is also used for ethical purposes, the majority of traffic on the dark web is explicitly illegal.

Dark Web Forum – A dark web forum is a website, accessible only via networks such as Tor and I2P, where users post and participate in discussion threads. Forums can cover a broad range of topics or be focused on a specific niche. Popular dark web forum discussion topics include the sale and use of drugs, carding and fraud, and hacking.

Dark Web Marketplace – Derived from “black market”, a dark web market is a site only accessible via anonymity networks such as Tor and I2P, that offers products and services in exchange for payment. The vast majority of dark web markets focus on the sale and purchase of illicit products and services, ranging from drugs, to hacking services, to weapons. Payments are typically made using cryptocurrencies, due to the increased anonymity associated with their use compared to traditional e-commerce transactions.

Data Leak – The release of data, giving unauthorized parties access to confidential information. This could be the result of a cyberattack, data deliberately stolen by a malicious insider, or an accidental leak due to human error. The information in data leaks can be used to execute further attacks.

Deep Web – The deep web refers to parts of the internet not accessible using standard search engines. It is made up of pages that are not indexed, including fee-for-service sites (like Netflix and Spotify), private databases (such as Dropbox and Paypal), and portals (for example, for Universities or organizations). It accounts for roughly 80 percent of the internet.

Distributed-Denial-of-Service (DDoS) Attacks

– A DDoS attack makes a service unavailable by flooding it with requests, to the point where legitimate requests cannot be dealt with.

Dox Sites – Doxing is the practice of finding and sharing personal or identifying information about a person or organization on the internet, usually with malicious intent. Dox sites are websites specifically dedicated to sharing these details, which could include their full name, home address, or phone number.

Escrow – A type of dark web marketplace where multiple sellers sell to buyers through one site. Buyers transfer their money to the marketplace, who holds the funds until the seller provides the goods or services.

Exit scam – An exit scam is when a dark web market's administrators withdraw the funds they are holding on behalf of the buyers and sellers and disappear.

Freenet – A decentralized dark web network that uses peer-to-peer connections to avoid censorship. Freenet is mostly used for file distribution such as blogs and publishing data that would otherwise be removed from the clear web.

Fullz – Dark web slang for a “full” packet of financial and personal information needed to commit fraud against an individual. “Fullz” are commonly sold on dark web marketplaces.

Incident Response (IR) – The process of containing and recovering from a data breach or cyber attack, with the objective of mitigating the damage to the organization.

Indicators of Compromise (IOC) – Forensic data, such as data found in system log entries or files, that identifies malicious activity on a system or network. Indicators of compromise aid information security and IT professionals in detecting data breaches, malware infections, or other threat activity.

Infostealers – Short for information stealing malware, infostealers are designed to extract valuable victim data that could be used to conduct fraud or other cyberattacks.

Internet Protocol (IP) Address – Every device connected to the internet is assigned a unique numerical identifier, an IP address, created by internet service providers (ISPs). Multiple users of the same network have the same IP address and it can be used to approximate a user's location.

Invisible Internet Project (I2P) – The Invisible Internet Project (I2P) is an anonymous network layer designed to facilitate private communication between its users. As it intentionally obfuscates activity, it fits the definition of what is commonly referred to as a dark web network.

Malvertising – Advertisements containing hidden programs that are used as a delivery method for malware.

MITRE ATT&CK – The MITRE ATT&CK framework is a free tool for organizations to map their defenses against the Tactics, Techniques and Procedures (TTPs) cybercriminals.

Multi-Factor Authentication (MFA) – An extra layer of protection used to ensure the security of online accounts beyond just a username and password. Can take the form of an authentication app, token, text message or other such additional security feature. Also known as Two-Factor Authentication (2FA).

Onion – Websites on the Tor dark web network are known as onions or hidden services. Domains on Tor end in .onion, preceded by a random combination of characters.

Open Source Intelligence (OSINT) – Data gathered from publicly available sites such as social media, government websites, arrest records, and news outlets. This intelligence is collected and disseminated for specific intelligence requirements.

Operational Security (OPSEC) – Originally a military term, OPSEC describes the process of identifying and obscuring information that could be gathered and exploited by an adversary.

Paste Bins – Text repository sites usually hosted on the clear web. These often contain items such as URLs or credential combos, as well as advertisements of data for sale.

Patch – An update to a program or its code that fixes a bug or vulnerability.

Penetration Testing (Pentest) – Simulated attack exercises, usually carried out by an external consultant, to identify security vulnerabilities and inform defense.

Phishing – A social engineering attack over email, text message, or direct message, which tricks the victim either to enter their credentials or to download malware via a malicious attachment.

Pre-Attack Intelligence – Threat intelligence that relates to the pre-attack tactics of threat actors, i.e. the actions they take before they breach the network of an organization. These tactics are defined in the MITRE ATT&CK framework as Reconnaissance and Resource Development.

Pretty Good Privacy (PGP) – Pretty Good Privacy (PGP) refers to an encryption program used for communication and authentication purposes. A public PGP key is used to encrypt a message before sending and the corresponding private key is then used to decrypt the message for its intended recipient. As each public key is unique, dark web users, especially market or forum admins, often use PGP keys to verify their identity and assure fellow users that their account has not been compromised.

Proxy – A proxy IP address is an online server that accepts and forwards requests for different devices on the internet. Hiding the user's true IP address, they are often used by threat actors trying to blend in with normal traffic while conducting malicious activity like data exfiltration.

Ransomware – A type of malware designed to deny a user or organization access to files on their computer by encrypting them. The malicious actors then demand a ransom payment in exchange for the decryption key, often also threatening to publish or sell potentially damaging information online if victims do not pay.

Ransomware-as-a-Service (RaaS) – A subscription-based ransomware model that enables affiliates to use already-existing ransomware tools to carry out attacks, while the RaaS operators take a percentage of every ransom payment.

Reconnaissance – The pre-attack techniques of cybercriminals, as they gather information on their target, before executing their attacks. This may include details of the victim organization, infrastructure, or staff/personnel.

Resource Development – The pre-attack techniques of cybercriminals, as they gather the resources they need to execute their attack on a victim. This includes creating, purchasing, or compromising/stealing resources that can be used to support targeting.

The Onion Router (Tor) – The Onion Router (Tor) is a dark web network accessed through free browser software. It enables more anonymous use of the internet and also hosts sites (known as onions) that are inaccessible through standard web browsers. Tor achieves this anonymization via a process called onion routing through thousands of relays – also called nodes – all over the world.

Threat Intelligence – Gathering information and data on cyber threats and threat actors to help inform defenses, to mitigate or prevent potential cyberattacks.

Threat Modeling – A process for identifying and prioritizing potential threats, based on threat intelligence, so that countermeasures can be developed.

Vulnerabilities and Exposure Management

Attack Surface – All of the assets, accounts, services, and exposures that could be discovered and targeted by an attacker.

Attack Surface Management (ASM) – The continuous discovery, inventorying, and monitoring of internet-facing assets and exposures to reduce the organization's visible attack surface.

Advanced Persistent Threat Group (APT)
– Groups that carry out attacks on a nation's cybersecurity or economic assets through acts of cyberespionage or sabotage.

Asset Inventory – A current record of systems, applications, identities, and cloud resources that need to be protected and managed.

Asset Discovery – The process of identifying systems, services, applications, and identities that should be included in security management.

Asset Criticality – The relative importance of an asset to operations, revenue, compliance, or business continuity.

Chaining – Combining multiple weaknesses, misconfigurations, or exposures to create a viable attack route.

Continuous Control Monitoring – The ongoing validation that security controls are deployed, functioning, and enforced across the environment.

Continuous Threat Exposure Management (CTEM) – A proactive cybersecurity framework by Gartner that continuously identifies, evaluates, and prioritizes security risks based on real-world exploitability and business impact. It shifts security from a reactive, tool-heavy process to a continuous cycle of risk reduction.

Control Gap – A place where a security control is missing, misconfigured, ineffective, or not being enforced as intended.

Compensating Control – A safeguard that reduces risk when the primary control cannot be implemented fully or immediately.

CVE (Common Vulnerabilities and Exposures)

– A standardized identifier for a publicly known vulnerability, used to track and communicate specific security issues.

CVSS (Common Vulnerability Scoring System)

– A scoring system used to rate the severity of vulnerabilities, typically on a scale from 0 to 10.

Cyber Risk Quantification (CRQ) – The practice of translating cyber exposure into financial or business impact to support decision-making.

Exposure – Any weakness, misconfiguration, or externally reachable condition that increases the likelihood of compromise.

Exposure Management – The ongoing process of identifying, prioritizing, and reducing the most meaningful exposures across people, process, and technology.

Exposure Window – The period of time between a weakness existing and it being remediated.

Exploit – Code or a technique used to take advantage of a vulnerability.

Exploitability – The likelihood that a vulnerability can actually be used by an attacker in practice.

External Attack Surface – The collection of internet-facing assets and exposures that can be discovered without internal access.

False Negative – A real issue that a tool or process fails to detect.

False Positive – A finding that appears to be a risk but is later determined not to be exploitable or relevant.

Initial Access – The first foothold an attacker gains inside an environment.

Internet-Facing Asset – Any system, service, or application exposed to the public internet and therefore potentially reachable by attackers.

Lateral Movement – Moving from one compromised system or account to another within an environment.

Mean Time to Remediate (MTTR) – The average time taken to fix vulnerabilities or exposures once they have been identified.

Misconfiguration – A security-relevant setup error, such as an open port, public bucket, or overly permissive access rule.

N-day – A known vulnerability for which a patch already exists, but many systems remain unpatched.

Patch Management – The process of identifying, testing, approving, and deploying software updates that address vulnerabilities.

Patch Latency – The delay between a patch becoming available and being applied in the environment.

PEM (Preemptive Exposure Management) – A proactive cybersecurity approach coined by Gartner, focused on finding, validating, and reducing exposures before they can be exploited. It emphasizes continuous visibility, prioritization, and remediation to shrink the attack surface as early as possible.

PoC (Proof of Concept) – A demonstration showing that a vulnerability can be exploited.

Prioritization – Ranking findings so teams fix the risks that matter most first.

PTM (Preemptive Threat Exposure Management) – PTM combines continuous attack surface management with real-world threat intelligence to close the exposure window before attackers can exploit it. Unlike traditional vulnerability management, PTM uses relevant threat intelligence and observed attacker intent to prioritize and remediate the most urgent exposures and threats.

Reachability – Whether an attacker can access a vulnerable service, system, or path from their starting point.

Remediation – The action taken to eliminate or reduce a vulnerability, misconfiguration, or exposure.

Remediation Backlog – The queue of vulnerabilities, exposures, or findings waiting to be fixed.

Remediation Workflow – The operational process for assigning, tracking, fixing, and validating exposure reduction tasks.

Risk-Based Remediation – Fixing issues according to exploit likelihood, asset value, and business impact rather than severity alone.

Risk Prioritization – The process of ranking issues so remediation effort goes to the exposures most likely to cause harm.

Shadow Exposure – Exposure that exists in trusted, widely deployed software, vendor products, or SaaS applications but is not fully visible or well understood by security teams. Unlike shadow IT, the asset itself may be known, but the real risk lies in the hidden vulnerabilities, configurations, dependencies, or attack paths it introduces.

Shadow IT – Technology, apps, or cloud services used without formal approval or visibility from security or IT teams.

Threat Modeling – A process for identifying and prioritizing potential threats, based on threat intelligence, so that countermeasures can be developed.

Validation – The process of confirming whether a vulnerability or exposure is genuinely exploitable.

Verification Scan – A follow-up scan or check used to confirm that a vulnerability or exposure has actually been resolved.

Vulnerability – A weakness in software, hardware, configuration, or process that could be exploited by an attacker.

Vulnerability Management – The ongoing cycle of identifying, assessing, prioritizing, remediating, and verifying vulnerabilities.

Vulnerability Scanning – The use of automated tools to find known weaknesses across systems, applications, and networks.

Zero-Day – A vulnerability that hasn't been either identified or addressed by those responsible for maintaining the technology, which can be abused by malicious actors until a patch is released. The term “zero day” refers to how long developers have to fix the issue before it becomes an issue for their customers.

The Preemptive Cybersecurity *Handbook*

Cybersecurity has spent decades getting better at detecting and responding to threats. But attackers are moving faster than defenders can react.

The Preemptive Cybersecurity Handbook outlines a different approach: continuously understanding your exposure, seeing your organization through an attacker's eyes, and acting before threats become incidents.

This book shows you how to put that approach into practice, with practical frameworks for changing how security decisions are made, how teams work together, and how success is measured.

It will help you build the capabilities, processes, and culture needed to move from reacting to attacks to getting ahead of them.

The result is a preemptive organization designed to disappoint attackers at every step.



SLCYBER . IO