Filigran

eBOOK

The Continuous Advantage

Why CTEM is the new imperative for AI-driven security teams

Table of Content

Introduction

CTEM breaks operational silos and translates cyber risk into business risk and management action

A lot of attention in cyber security is directed to unknown external forces: sophisticated attackers, tricky techniques and exotic hacking tools. But that's not the biggest challenge for many organizations. Instead, it's knowing your defenses, and knowing if they're up to the job.
The key questions: What can hurt our organization? How would that happen? And finally: Can we stop it from happening?

These questions aren't just cybersecurity questions: they're core considerations for leaders throughout an organization.

In this eBook, we're going to explain how Continuous Threat Exposure Management (CTEM), as a framework, helps organizations answer the security questions that matter most - **what is exposed, how likely is it to be exploited, and where should we act first**, in an environment where insights are often scattered, incomplete, and hard to act on.

According to Gartner, 94% of the vulnerabilities out there are never exploited. Of the remaining 6%, a further small fraction are a concern to your organization. There are plenty of tools and capabilities identify all of those but it can get complicated and confusing very quickly, even for the most capable teams.

And now with frontier AI and recent initiatives like Project Glasswing, we are realizing that as discovery of vulnerabilities leapfrogs, security teams will find it increasingly difficult to pick up and fix the ones that truly matter to them, unless the loop between discovery and validation is closed.

Understanding how to harness this to prioritize vulnerability management, and how to update that prioritization as both the threat landscape and the organization changes around it, is the problem that CTEM seeks to address.

There isn't a shiny, shrink-wrapped product that will answer this question: it's as much about people, culture and organizational setup as anything else.

The technical teams needn't despair, however: with the right threat intelligence, adversarial validation and crisis simulations tools, CTEM becomes a continuous feedback loop that absorbs threat intel and validates and prioritizes in a near-constant process.

A Definition: Continuous Threat Exposure Management

Continuous Threat Exposure Management (CTEM) is an iterative framework that helps organizations identify, assess and remediate cyber security weaknesses across their digital estates. CTEM's five-stage iteration replaces traditional periodic assessments with a way of continually understanding and managing your cyber risk.

Gartner introduced CTEM in 2022 to rectify shortcomings in traditional vulnerability management practices. Traditional vulnerability management focuses on generic vulnerability data and scores but omits to cater to the far-from-generic nature of most organizational risk profiles.
A second issue is that periodic assessments are static, and often a reaction to a regulatory requirement. Things like pentesting are only as accurate as the timespan in which they are run, providing a mandated snapshot of security readiness that is outdated within hours or days.

Gartner defines CTEM as, "A pragmatic and systemic approach to continuously refine priorities for remediation efforts and trigger change in security posture by assessing the accessibility, exposure, and exploitability of assets."

CTEM combines and improves upon these two issues, going beyond mass-digestion of vulnerabilities to look at context: how an organization is exposed, and which of these potential exposures can be exploited. Validation tools are then employed on a more continuous, rolling basis.
Where traditional assessments were rigid, periodic and often ignored significant and continuous change, **CTEM embraces a continuous, iterative cycle through five distinct phases or stages.** Once the final stage is reached, work starts again on the first: Scoping. We'll explain more about this in detail later, but for now, let's look at the impact on businesses.

Why should business leadership care about CTEM?

CTEM brings together cyber risk and business risk in a language and a framework that is easily digested by both technical and business people. It helps organizations build relevant and accurate, outcome-focused risk scores that help security teams prioritize remediation efforts, and business leaders understand the business risks (and opportunities) effective cyber defense provides. In turn, this understanding helps inform budgeting and make the case for security investments. Two of the key deliverables of a CTEM program are proof of this: alongside a machine-readable JSON feed, CTEM should also deliver documents in plain language for any human to be able to read.

CTEM builds a dynamic picture of exposure and cyber risk - one that evolves alongside both the organization and the threat landscape. Unlike static assessments, **it is continuous and iterative by design**. Cyber threats shift constantly, accelerated by the growing use of AI on both sides of the equation, by attackers and defenders. As a result, what is relevant and critical today may not be tomorrow. CTEM addresses this directly. **It equips security teams with the ability to continuously test and assess their readiness against the threats that actually matter to their environment - cutting through noise rather than adding to it.**

What does CTEM actually do for an organization?

We've touched on the ability of CTEM to help both technical cyber teams and the wider business. A lot of this value comes from using CTEM to identify what could go wrong, how it might happen and what the organization can do to minimize or eliminate damage.

Arguably the most valuable outcome is that using this operating model describes cyber risk in language the rest of the business can understand and use.

Here's a run-down:

ATTACK PATH MAPPING	Firstly, CTEM helps identify the potential attack pathways that directly affect the organization.
GOVERNANCE – WITH EVIDENCE	Next, it provides evidence-based governance of these threats. In plain English: how to manage and minimize the risk these threats represent, and a reduction in the potential fallout from attacks.
FROM FIREFIGHTING TO RISK ANALYSIS	An early output is shifting teams from a response-led approach, reacting to incidents and attacks as they occur, to an intelligence-led, proactive security approach that allows the organization to switch to viewing cyber risk as a business risk.
INSIGHTS YOU CAN USE	CTEM then also collates potential remedies to problems that already or will affect the organization, as opposed to listing a load of problems that may or may not be relevant. An example: risk and compliance leads within organizations can map CTEM categories to Key Risk Indicators (KRIs), Cyber Risk Quantification (CRQ) and Key Performance Identifiers (KPIs) to quantify trends over time. This can come especially handy for existing security tools' renewal time decisions

A VALIDATION OF PEOPLE AND PROCESS, AS WELL AS TECHNOLOGY

Although, CTEM as an operating model, focuses on security tools' validation but we would say that security posture is a combination of tools, processes and people, and therefore we do need to focus on overall state of the organization's cyber security defenses, not just its technological components.

A SOFT NUDGE TO HELP BREAK DOWN SILOS

Another reason your leadership team will like CTEM: it requires engagement and information sharing from multiple teams. CTEM should never happen in a silo, and never just involve the security team.

IT'S NOT A SHINY OBJECT FOISTED ON EMPLOYEES BY CONSULTANTS

At least, it shouldn't be. CTEM is not a shrink wrapped piece of slideware, and nor is it a cure-all. It's for organizations to build and own themselves as an operating model, and that does require some change management, but it also provides an opportunity to do a little silo breaking in the process.

CTEM needs senior management buy-in for success

By now, it should be clear that **CTEM is not a model** to be applied by operational management, or even by the CISO or CIO alone: **it requires buy-in from multiple teams, and that calls for heavyweight support.**
The good news is that for most organizations, tools that help improve an understanding of risk and risk postures are welcome. **CTEM helps the financial, compliance and governance teams and the senior business leadership, as well as the security function itself.**

The tricky part? Sometimes vanity metrics are well-established and jealously guarded, especially if they show that everything's fine (well, until it's very much not fine, that is). The argument for evidence-based security outcomes is always worth making, but framing it in such a way as to avoid calling out vanity metrics for what they are can be a challenge.

The five stage CTEM cycle

So, your organization is poised to adopt CTEM as an operating model. The CEO, C-Suite and senior managers are all active supporters and promoters. It's time to start the first CTEM cycle.

We're going to spend more time on the first step, Scoping, since it is the most critical point in the initial CTEM cycle. Get it right, and the various teams involved are working with a pragmatic, risk-centered scope that is achievable and measurable. Get it wrong, and they're working with a scattered or broad, unmanageable effort.

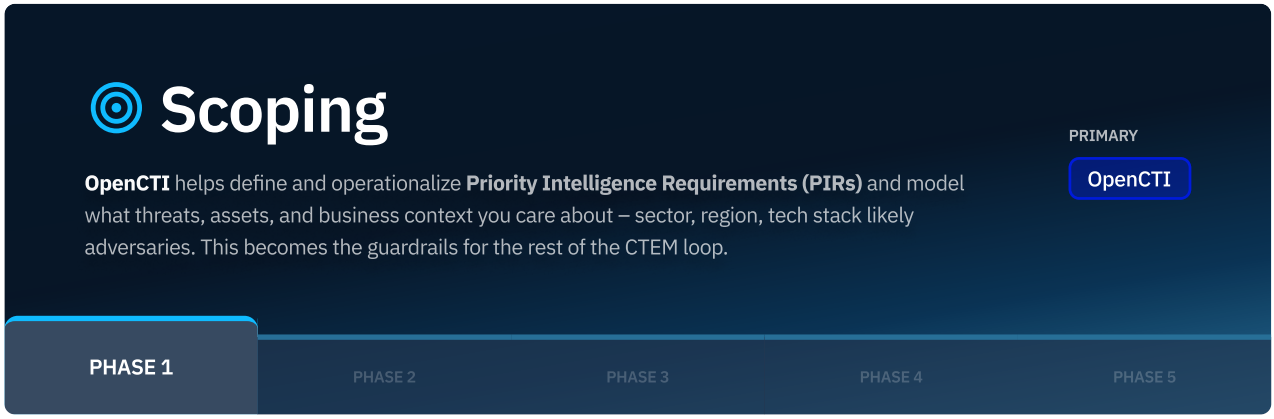

PHASE 1: SCOPING

The first phase requires that buy-in we've discussed, because it looks at what makes up the critical assets, capabilities and processes of the organization. This stage is effectively the 'what and where' part of the question **CTEM answers: what needs to be included (to be protected), and where to focus first.**

Look first for the business risks that effective cybersecurity can directly influence. Secondly, work out which environments are in scope for this iteration. The scope can be changed or expanded in future go-rounds, but for now it's important to make the first such cycle achievable. Environments in this context are either technical, things like on-premises or cloud computing, IT/OT systems or organizational, such as subsidiary businesses.

The same applies to asset types. **Should you include all endpoints? Can you realistically do so? What are the company or organization's crown jewels? Finally: is your view of this inventory accurate and up to date?**

Bear in mind that all of this activity and planning should take into account what resources can be pointed at problems as they're identified, so a pragmatic approach is worthwhile. Planning should include the capacity of the security team, existing service level agreements from outside suppliers, and the tooling available to do the job at hand, as well as any operational budget that might need to be pulled in.

What role does OpenCTI play in all of this?

OpenCTI

OpenCTI is a powerful threat intelligence platform (TIP) for sorting through the mass of potential vulnerabilities from multiple intelligence feeds to find the relevant risks for an organization.
OpenCTI is a powerful assistant in the **first three stages of the CTEM cycle**, developing Priority Intelligence Requirements (PIRs) for the scope, delivering threat-led prioritization to OpenAEV during the prioritization stage and applying structure and making connections with STIX 2.1 and MITRE ATT&CK mapping to help explain exposure during the Discovery phase. It is also invaluable in helping maintain the continuous intelligence loop that continuous reassessment under CTEM requires.

Another way of looking at the Scoping phase is as a more in-depth process of building Priority Intelligence Requirements (PIRs); Cyber Threat Intelligence (CTI) platforms like Filigran's OpenCTI can help to define and operationalize PIRs, modelling the threats, assets and business context.

The next step is to find and fold in assessments of threat actors and attack methods relevant to your organization's sector and tech stack. Look at how to use this insight to refine the scope. You'll then need to define what critical exposure looks like for your organization. The impact of data theft, denial or exposure is probably going to be the most obvious measurement, but data sensitivity, ease of exploitation and likely impact on customers, supply chains and other groups should also be anticipated. Timelines for detection, response and recovery will also have a bearing on this measurement.

Finally the teams must figure out the People, Process and Technology they currently use, and what capacity they have to fix the problems identified in the Scope. This includes key elements, such as Service Level Agreements (SLAs) from service providers, vendors and Incident Response specialists. **But it should also include an honest assessment of the capabilities of the tools and people in play, too: do the skills and resources exist to make the difference when it matters?**

This is not an exhaustive list, but these questions help define a realistic, risk-aligned CTEM scope that can be executed and measured, instead of an overly broad, and therefore unmanageable, effort.

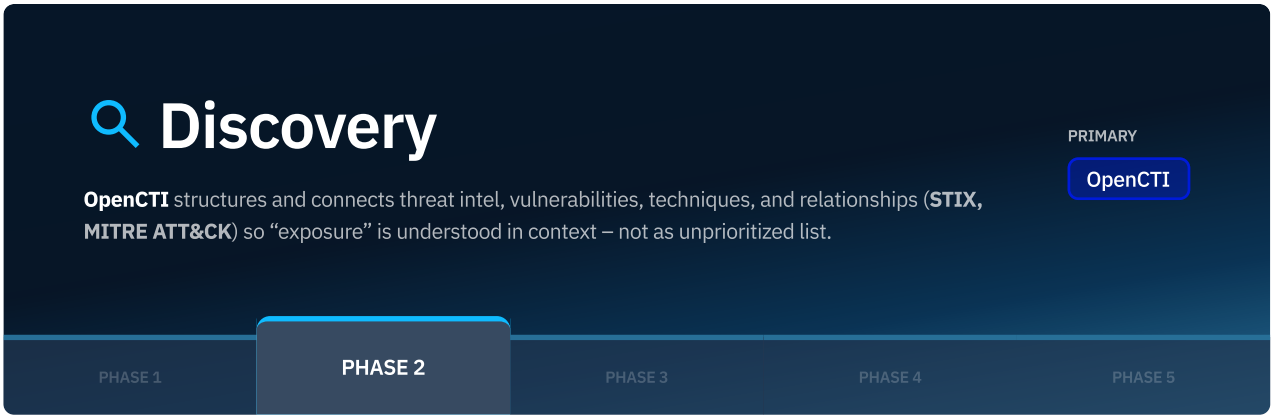

PHASE 2: DISCOVERY

The second stage of the CTEM cycle is to take the data from the Scoping phase and build a comprehensive picture of what is exposed within those boundaries, taking the initial list of what matters to the organization and putting it in context for the third stage: Prioritization.

Within those defined boundaries of scope, it should now include complete picture of the attack surface; everything from unmanaged devices to air-gapped OT environments and Cloud services, and all the stuff in between: IT infrastructure, including Cloud and SaaS, IoT and unmanaged devices. Identifying and eliminating blind spots is a crucial part of this step. Look for unmanaged devices, as well as Shadow IT and Shadow AI brought in by employees without permission. A modern Attack Surface Management (ASM) tool should be able to provide visibility into this.

A few examples of **what to look for:** outdated Operating Systems or applications, default credentials in use, insecure protocol usage, misconfigurations and poor network and domain segmentation, and CVE (Common Vulnerabilities and Exposures) matches. Also look for external-facing assets suffering from issues with security controls.

With the continuing impact of supply chain attacks such as those on Asahi, Jaguar Land Rover and Marks & Spencer, it's also a timely reminder to understand the full scope of both up - and downstream supply chain dependencies and the risk they may represent, not just to the high-profile victims, but their suppliers and customers.

Look also at less tangible assets. Things like the organization's social media accounts, code repos on third party sites like GitHub, and integrated supply chain systems, for example.

Finally, take an external perspective: what do all of these assets look like from the outside to attackers? This is where ASM helps, but also where Cyber Threat Intelligence (CTI) tools like OpenCTI can make threat-led prioritization a smoother process, while Adversarial Exposure Validation (AEV) tools such as OpenAEV turn that output into test scenarios that can be mapped to specific Tactics, Techniques and Procedures (TTPs). Modern Exposure Assessment Platforms (EAPs) can also be put to good use here to pull together data from multiple sources. Standardization frameworks such as MITRE ATT&CK and STIX 2.1 are helpful too, to build a consistent view of threats and possible attack paths.

PHASE 3: PRIORITIZATION

Prioritization is where effective management of threat intelligence comes into play. The aim is to move from the reactive behavior we described earlier, to more systematic and pre-emptive strategy that closes the exposure window, identifying immediate and real-world risks. Prioritization also lays the foundation for evidence-based governance (critical to the Validation phase) and helps build operational resilience. From a technical and operational security standpoint, the evidence collated and presented at this stage also makes it easier to demonstrate the need for any urgent remediation spend or activity at the fifth stage – mobilization.

Remember that 6% of exploited CVEs we talked about? This is where the data and insight from the discovery phase is coupled with Adversarial Exposure Validation, starting with the missing 6% of CVEs that matter, and thin this list out even further.

A key task is to understand which potential adversaries and TTPs are most relevant. This is where PIRs come into play, defining threats based on sector, region, tech stack and organizational context and as a result alerting/ prioritizing threats that really matter and are important to be neutralized. This can then be measured against the business assessment of asset criticality from the first stage of the CTEM process.

Finally, the results should be put in context with the available controls: Endpoint Protection (EPP) or Endpoint Detection and Response (EDR) are good examples, equipping organizations with the ability to fine-tune the security environments for endpoints according to their exposure.

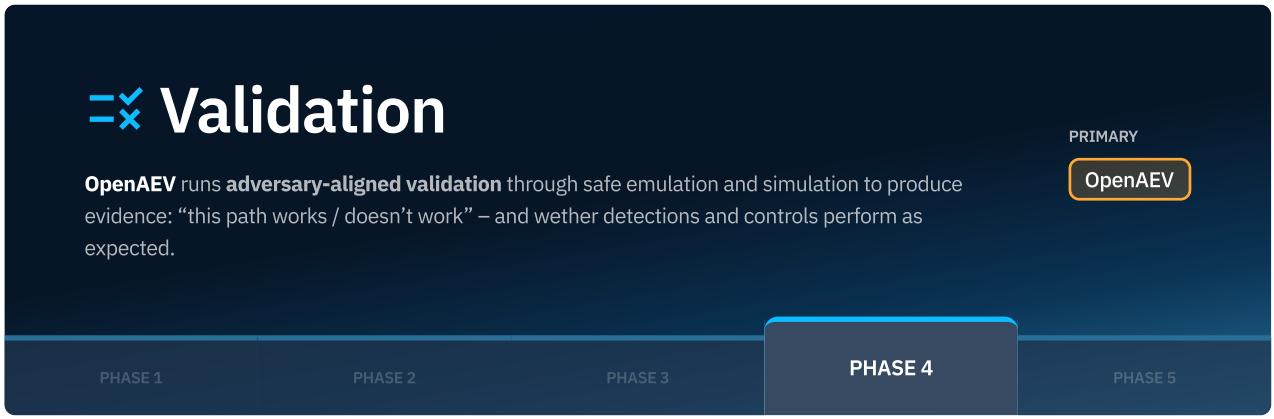

PHASE 4: VALIDATION

This is the point at which all of the discovery, analysis and theory-building is tested. The Validation stage answers the question at the heart of a CTEM assessment: **Can we stop attacks?**

Validation is also the point where the security team can start to run controlled attack emulations to test the validity of the risk scoring and vulnerability modelling that has taken place in the previous rounds.

The attack emulation process should evaluate the actual exploitability of the potential vulnerabilities identified in the first steps of the process.

Next, it should map the attack pathways – the potential routes an attack could take to reach the organization’s valuable or sensitive assets.

Finally, it needs to demonstrate control efficacy: do existing tools and controls, such as EDR or the organization’s SOC and SIEM setup, actually detect and block the synthetic threats being tested. And, take remediation action to fix the gaps.

Nowadays, this process is becoming more automated and continuous, thanks to AEVs like Filigran’s OpenAEV. Adversarial Exposure Validation combines automated penetration testing, breach and attack simulation (BAS) activity and cyber range exercises into a continuous capability that removes a lot of the heavy lifting from teams and reduces the gaps between what have historically been quite labor-intensive activities. An outcome of this (and, arguably, a cornerstone of CTEM itself) is the creation of evidence-based and outcome-focused governance.

It would be criminal not to mention the People, Process and Technology triangle at this point. As is now hopefully very clear, leaning too much on the technological leg of the tripod creates a weak, reactive security posture.

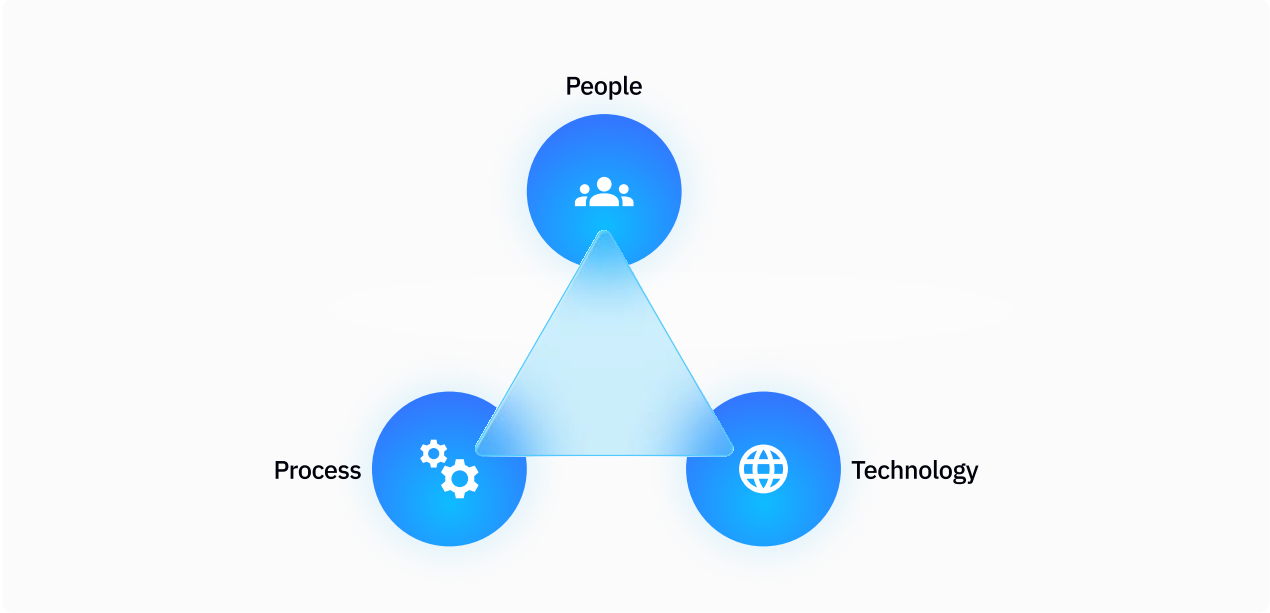

As usually described, the Validation stage of CTEM is focused almost entirely on tooling, and the process of checking that those tools are up to the task at hand. Any technical testing to validate security tools must also come with table top and crisis management exercises to exercise the processes and people involved. This more thorough validation is something OpenAEV was built to do, taking it beyond traditional AEV capabilities.

An effective use of CTEM should set out to right this imbalance, validating processes, testing the readiness of teams and not just the security team. The result casts the organization's defenses as a single, integrated system that is adjusted and tuned to meet changes in its environment.

What does AEV do for organizations?

OpenAEV

OpenAEV is incredibly helpful in the CTEM process at the points of validation and mobilization, but also at the final ' GOTO: 10' stage of going back to the start of the CTEM cycle.
During the Validation process, OpenAEV runs adversary-aligned validation through safe emulation and simulation to produce evidence that can be used in the Mobilization stage, where it creates usable findings and helps move workflows to task owners' inboxes.
Finally, we keep talking about People, Process and Technology, and OpenAEV validates and tests this trinity with the ability to execute both technical exercises and generate realistic tabletop simulations.

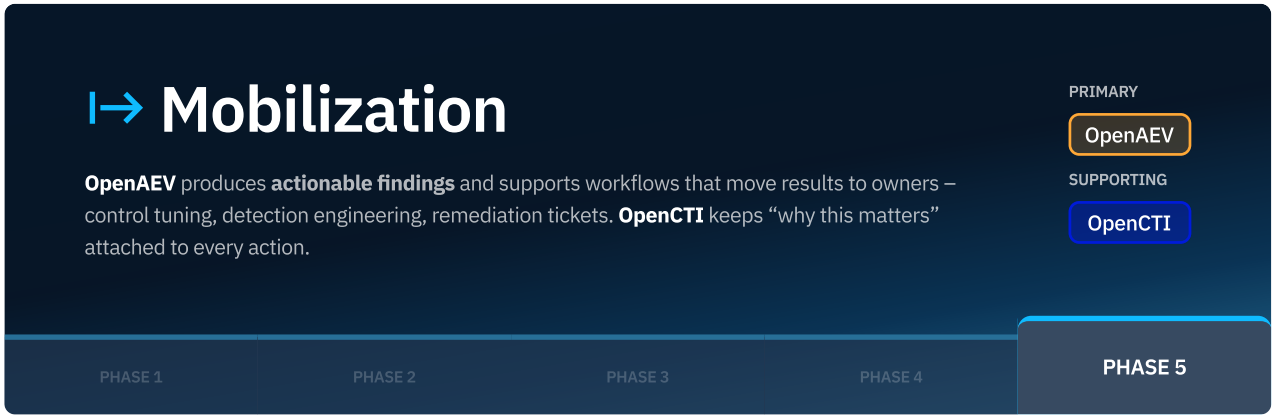

PHASE 5: MOBILIZATION

The last step before going back to the start for a fresh lap is to take what has been learned in the four previous steps and turn it into the evidence needed to justify change, remediation and eventually the correct risk calculations for the organization.

This is the point during the process when cyber, information and risk teams can start to operationalize the learnings of the previous stems, and of remediation itself. In Plain English: it's the delivery of a To Do list to each relevant team, usually via incorporation into a ticketing platform and, increasingly, the use of AI to automate the assignment of tasks using inference and predefined rules to pick the right group, team or individual for the job. Automated responses can, and likely will, extend to actual interventions to isolate or neutralize compromised endpoints or services.
Key to mobilization is a commitment to evidence-based action by the people involved. Security teams should, for example, take results from the Validation stage and use it to make the case for remediation requests. As the same people signing off on this should be wholehearted participants in the CTEM process, the outcome is hopefully a more efficient and speedy remediation process.

Validation tools like OpenAEV provide remediation guidance with multiple options, making it easier for customers to take the most prioritized remediation step first and then work through the list. Mobilization stage also helps with evolution and evaluation of security tools to better integrate and deliver the context of other parts of the organizations.

Starting all over again

By the time the organization reaches this point, the value of CTEM should be clear: clarity, measurable risk and measurable outcomes are three of the outputs. Repeating this process continues to deliver actual results, rather than anticipated savings. Organizations and threat landscapes change, and as risk appetites and goals also evolve over time, and an initial assessment still remains fixed in one point in the past. Reiteration of the process also allows organizations that are resource-poor to iteratively improve their security and risk posture, chipping away at the problem one step at a time.

Four key ingredients for a strong CTEM framework

Technology platforms

As part of the People, Process and Technology triangle, Technology is often painted as the shiny, distracting object. But it's still vital. Organizations should look to use CTI, AEV, CRQ and ASM platforms to understand and test their potential areas or vulnerability

Threat Intelligence

Equally, it's easy to become swamped with a high volume of threat intelligence, and find it impossible to track that 6% needle in the 94% haystack. Strong CTI capability and effective use of PIRs are key to avoiding overload.

What will rapid iteration and improving capability in LLMs / AI do to the CTEM framework?

In short: it's already having an impact on workloads, automating a lot of the repetitive tasks and evidence-based judgement calls that teams need to make during an assessment cycle. Reducing the workload and learning curve for teams adopting CTEM, and helping speed the validation and adaptation to new risk landscapes, can result in faster and easier iteration of the CTEM process.

AI and Automation

AI and automation is almost an essential component or rather enabler of CTEM as it won't manually be possible to run such a program on a continuous basis. AI-assisted automation accelerates and repeats operations that were previously manual and time-intensive including threat intelligence structuring and processing, penetration testing simulations, attack path analysis, and control validation. It creates speed and scale without bogging teams down in toolmaking and data sorting, and allows for continuing adaptation. Agentic AI can take on defined, repeatable tasks autonomously - triaging findings, assigning remediation actions to the right teams, and escalating what genuinely requires human judgment. This removes the coordination overhead that slows security programs down, particularly in organizations where responsibility is spread across multiple teams or business units.

Holistic Assessment

Returning to the People, Process, Technology angle, organizations benefit from winning and maintaining a complete picture of their environment from GitHub repos to smart lightbulbs in the warehouse and everything in between. Understanding attack pathways is a key benefit, but a richer understanding of the environment benefits any organization.

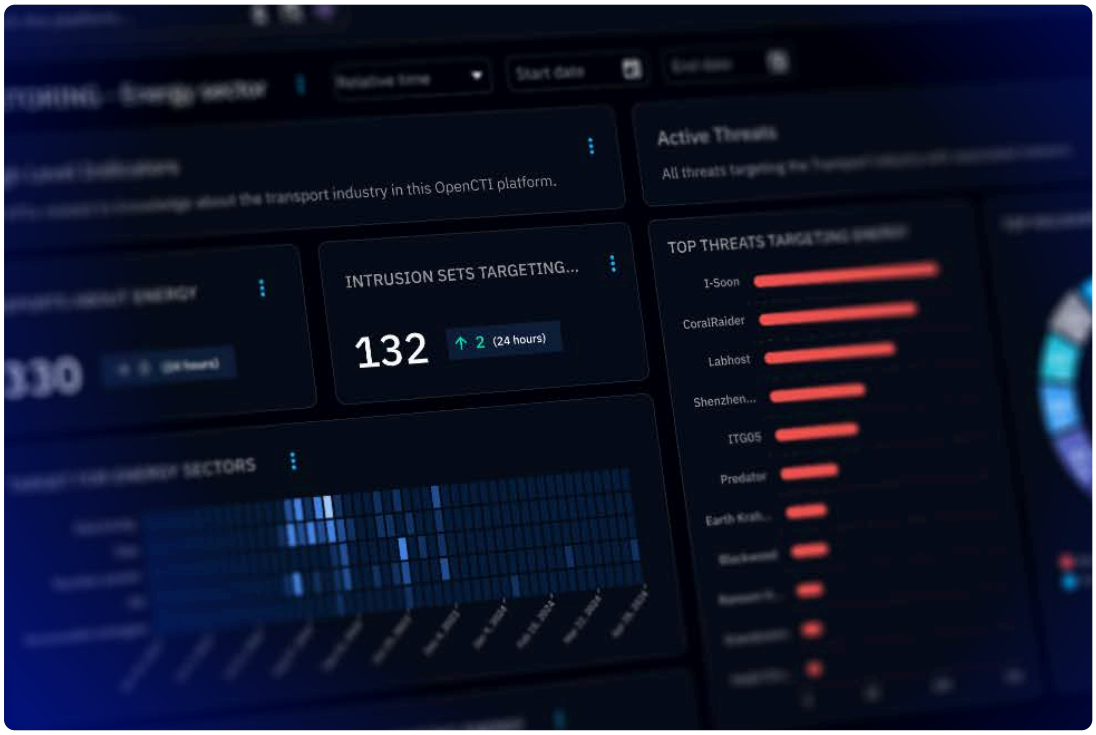

Why leadership and governance is the critical ingredient

CTEM only works as a leadership initiative; it's not a project to sling casually in front of a selection of teams with the instruction to make it happen: the organization's leadership cannot pass on the responsibility for making CTEM work. Organizational leaders, from the board and CEO down, have a pivotal role - and core responsibility - for making a CTEM initiative successful. The reward is a less siloed organization and one that is able to take managed cyber risks when lesser competitors must either back down or hope for the best.

Conclusion

For years there's been talk of the importance of CISOs and other tech leaders talking business risk to the Board. CTEM helps with this, for sure. But CTEM is much more than that: a clear operational framework to help cut through hype, noise and chaos to find (and constantly re-find), quantify and tackle the cyber risks they face.

Disconnected tools create disconnected programs. Discovery without intelligence context creates noise. Intelligence without validation creates anxiety.
Prove that your defenses actually hold against the threats that are actually coming for you.

We need to close the loop; where threat intelligence directly drives what gets validated, and validation results feed back into what gets prioritized. If that loop is not closed in your environment today, closing it is the highest-leverage investment you can make.

ABOUT FILIGRAN

Filigran stands out for its expertise in open-source solutions for end-to-end threat management. It offers the Filigran **eXtended Threat Management (XTM) suite**, designed to help organizations understand their threat environments, anticipate and detect incidents, and strengthen their overall security posture. The suite currently includes two solutions:

- **OpenCTI**: structures and operationalizes cyber threat intelligence at technical, operational, and strategic levels.
- **OpenAEV**: helps identify critical security gaps through advanced attack simulations, resilience testing, and crisis management exercises.

Today, more than 6,000 organizations worldwide rely on Filigran's solutions. With a team of over 200 employees across the globe, Filigran supports leading companies such as Marriott, Thales, Hermès, Airbus, Novartis, and Bouygues Telecom, as well as for public entities including the European Commission, ENISA, ANSSI, the French Ministry of the Interior, the New York State Cyber Command, and several U.S. and Australian federal agencies. Filigran has also built a strong ecosystem of partners, including Deloitte, Orange Cyberdefense, Deepwatch, Arctic Wolf, Google Cloud Security, Atos, Wavestone, and Intrinsec.