5 steps to build your CTEM program

Continuous Threat Exposure Management

A 5-Step Guide to Proactive Cyber Defense

Continuous Threat Exposure Management (CTEM)?

And, how is it different from traditional exposure management?

Continuous Threat Exposure Management (CTEM) is Gartner's framework for continuously identifying, validating, and reducing an organization's cyber exposure. Unlike traditional vulnerability management, which often produces long lists of technical flaws with little context, CTEM is a programmatic approach that aligns security efforts with business priorities and real-world threats.

How does Gartner define CTEM?

Gartner defines a CTEM program as "a set of processes and capabilities that allow enterprises to continually and consistently evaluate the accessibility, exposure and exploitability of an enterprise's digital and physical assets".

What is the goal of CTEM?

In practice, CTEM shifts organizations from a reactive, siloed mentality to a continuous, risk-informed cycle. Its key steps include scoping, discovery, prioritization, validation, and mobilization.

Why CTEM and why now?

In today's environment of expanding attack surfaces spanning cloud services, remote workforces, IoT devices, and third-party integrations, old methods are failing.

Organizations can't fix every potential issue nor confidently decide how to prioritize with a plan, leading many to operate with significant blind spots. Meanwhile, executives are pressing security leaders for assurance that cyber investments actually reduce risk to the business. This can't simply be achieved by endless technical vulnerability lists.

While this is not a move in the wrong direction and can still be proactive, it still leaves a lot to be desired. It doesn't consider the human-side of vulnerabilities and often can lack business-relevant prioritization.

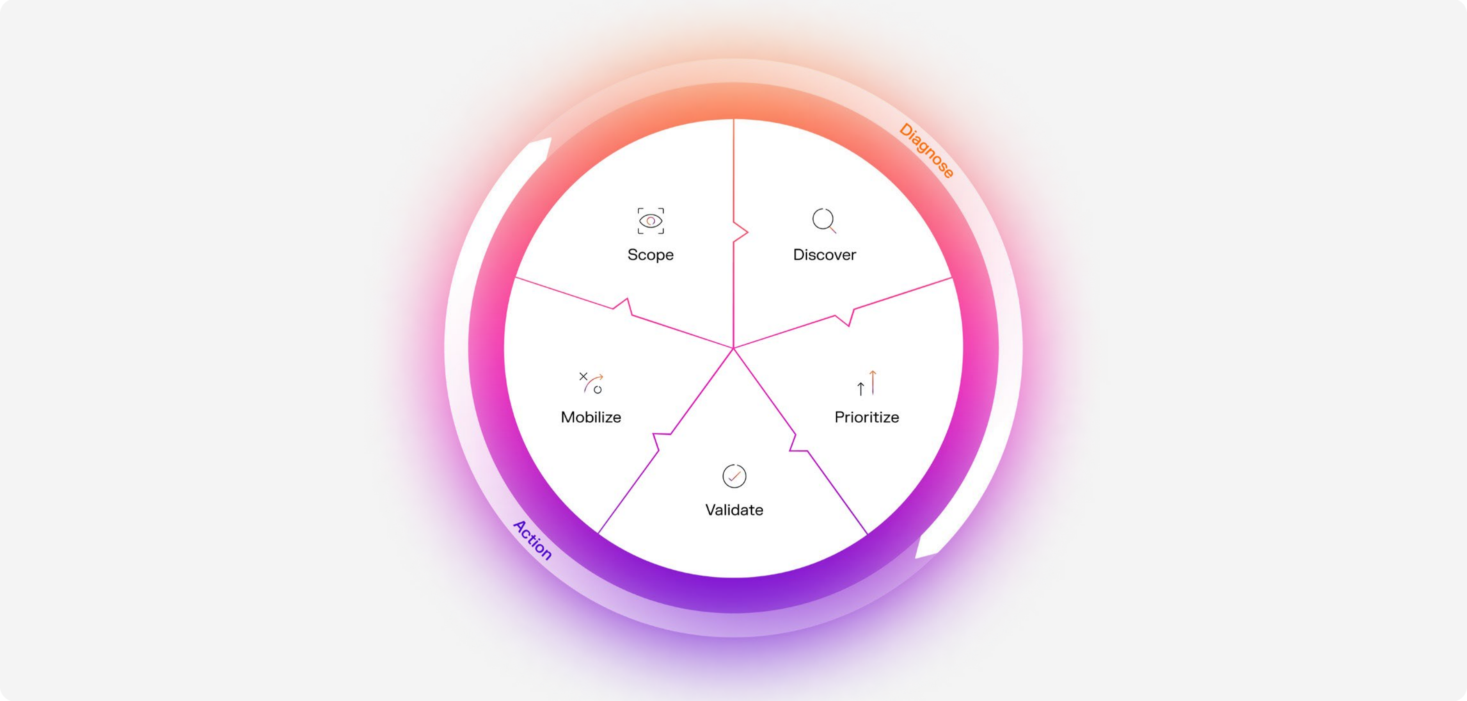

Figure 01: The CTEM 5-step approach workflow

Growing CTEM Adoption

Gartner predicts that by 2028, organizations that implement CTEM will see at least a 50% reduction in successful cyberattacks. However, CTEM is not yet universal. Many security teams struggle to embrace and implement CTEM due to several key challenges:

Siloed, tool-centric approaches

Vulnerability management, incident response, and threat intelligence teams typically use disconnected tools and processes, leading to fragmented data and teams wasting time triaging noise instead of addressing critical risks, resulting in important exposures falling through the cracks.

Focus on vulnerabilities over business impact

Traditional programs often prioritize based on severity scores (CVSS) or generic threat feeds alone, ignoring the actual business impact of an exposure. High-impact risks linked to critical assets may be underprioritized, while low-relevance issues consume resources.

Tunnel vision of risk factors

Periodic scanning misses dynamic cloud assets, SaaS activity, and human factors. A solely technology-focused view overlooks things like that can't be patched or can't be quantified at a high-level or generic lens.

Lack of cross-team mobilization

Even with a prioritized list, security teams often hit a wall trying to remediate issues that require action by IT, DevOps, or business units. Stakeholders need to be convinced and involved in order to fix issues without a massive stall.

Difficulty measuring value

Risk reduction is difficult to quantify in business terms. Many programs lack outcome-based metrics to demonstrate progress, making it harder to justify purchasing or maintaining products that support CTEM efforts to budget owners, especially when compared to more straightforward and defined compliance checklists.

Despite these challenges, the need for CTEM is clear. Attackers exploit whatever gaps they can find whether it's an unpatched server, a misconfigured cloud container, or stolen credentials from a third-party breach. CTEM provides a structured, continuous way to find and address these gaps before attackers do, and to do so in a manner that balances lean resources against the most critical threats to the business.

Source: Use Continuous Threat Exposure Management to Reduce Cyberattacks by:
Gartner - Jonathan Nunez, Pete Shoard, Mitchell Schneider

Whether an organization is a young startup with less than 100 staff or a 10,000-person multinational banking corporation, there is something for a business of any size and background to gain from properly embracing a CTEM framework into their security practice.

The following 5 steps address both what to do and why it matters, with real-world insights for overcoming adoption challenges.

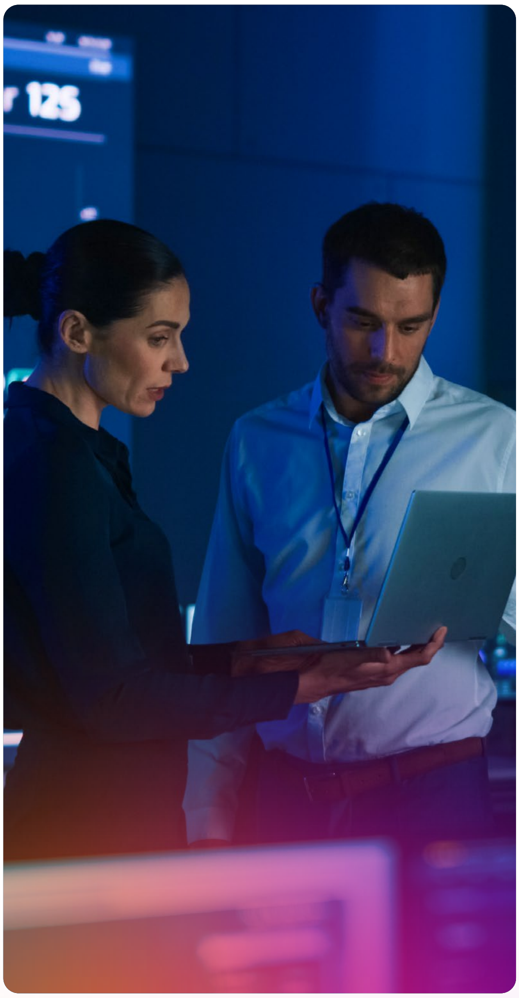

The 5 Steps to Proactive Cyber Defense

Step 01

Scope exposures based on business priorities

The first step in CTEM is to define the scope of your continuous exposure assessments in alignment with business impact.

Rather than indiscriminately scanning everything, effective CTEM programs target the areas that pose the greatest risk to the organization's critical objectives. It's important because security teams often struggle to get leadership support when they present technical findings without context.

By contrast, when exposures are framed in terms of business impact, executives and budget owners are more likely to allocate resources.

Foundational questions:

What are our organization's most critical business processes and assets? These can be revenue-driving customer portals, sensitive intellectual property or production systems.

Who owns these processes and assets? Identify the business owners and technical owners. These will be your primary business owners who will be involved in remediation.

Where are these assets and what are their dependencies? Map out the IT systems, cloud services, third-party applications, and data stores that underpin key business areas. Don't forget things outside your direct control: third-party SaaS, integrations, open-source components, social media accounts.

What worst-case scenarios are we worried about? This question helps scope scenarios most relevant to your organization's risk appetite. Everyone's definition of worst-case is different.

Exposure scopes

Answering these questions defines your exposure scope. Instead of a vague scope like "all Internet-facing assets," you might scope "exposures in our public-facing web applications that handle customer transactions". The idea is to group assets and exposures in a way that anyone instantly grasps the relevance. For organizations new to CTEM, pick an initial scope that is high impact but easier to manage.

A common best practice is to begin with the external attack surface of a critical service. External assets are often a priority, since they are exposed to everyone and are relatively easier to scope. Other popular starting places are SaaS and cloud services, as many firms have visibility gaps there.

Run multiple scopes in parallel

You do not have to tackle exposures one scope at a time. In fact, Gartner notes there is "no limit to the number of scopes that can be run concurrently."

Mature programs often operate multiple CTEM workstreams simultaneously, each addressing a different domain, with a "master scope" for reporting up to leadership and more granular technical sub-scopes underneath. Each scope should have its own cycle of discovery and remediation, but all feed into a unified CTEM program. By scoping in this way, you make the CTEM effort relevant to stakeholders. Business leaders don't need or want to see a list of 10,000 vulnerabilities across the enterprise; they want to know "How exposed is our online customer portal?" Scoping provides that focus.

By starting with exposure areas tied to business priorities, you create early "wins" that demonstrate the value of CTEM, building momentum for broader adoption.

Multiple CTEM Scopes running simultaneously

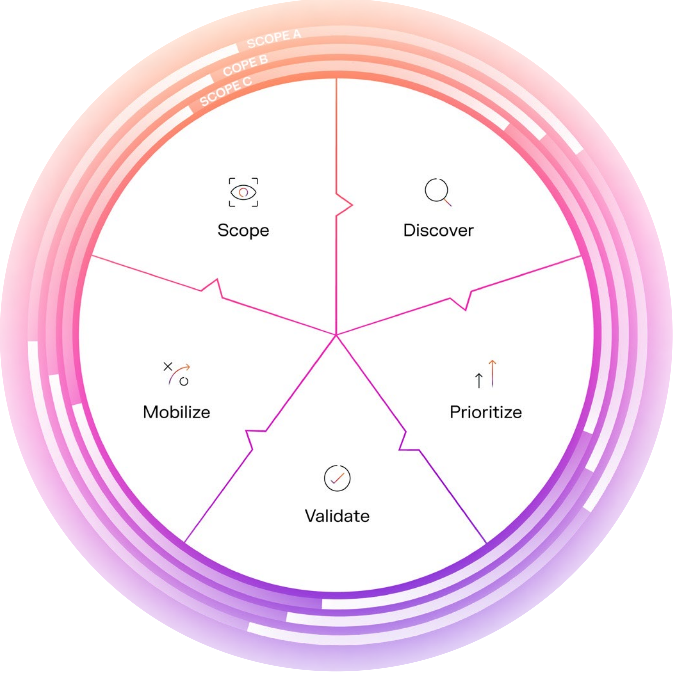

Figure 02: multiple ctem scopes as described by Gartner's Strategic Roadmap for Continuous Threat Exposure Management

Step 02

Validation-driven prioritization

Once you have scoped and begun discovering exposures, you will inevitably end up with lists of issues: vulnerabilities, misconfigurations, identity weaknesses, missing controls, etc.

The next step is to filter and rank these findings by importance (or use a tool is automates this task) and then validate the most critical ones through simulated attacks or testing.

Prioritization with business context

CTEM takes a more nuanced approach to vulnerability management compared to using solely letter grade or number ranked asset lists. Useful factors to consider include likelihood, criticality, accessibility, mitigating factors, and impact. By incorporating these, you perform risk-based prioritization. For instance, you might downgrade a critical vulnerability that has no known exploit and is on a segmented network but elevate a medium-severity misconfiguration that could allow an attacker already inside to pivot to your finance database.

Threat intelligence should feed into prioritization. If credible intel suggests a particular exploit is on the rise targeting your industry, exposures related to that should move up your list. Many modern risk-based vulnerability management tools incorporate such intelligence. However, don't rely solely on external scores.

Validation

Once you have a prioritized list of exposures for a given scope, take it one step further and validate the top risks. Validation means testing whether the exposure is truly exploitable and understanding what happens if it is exploited. In practice there are many ways to do this:

Adversarial simulation: Using breach and attack simulation tools or manual red teaming for the highest risks to safely attack via the pathway a real adversary would take. For instance, this could look like attack engagements of externally facing web applications or phishing simulations of high ranking or impactful individuals with greater digital access.

Security control assessments: Verifying that defensive controls would detect or prevent an attack leveraging the exposure.

Configuration tests: Sometimes exposures are just a config issue that can be tested. For instance, this can look like using default credentials on a service in a login to confirm the level of access.

Automated validation tools: Quickly check large numbers of findings for "exploit success" and subsequent pivot opportunities in sandboxed methods.

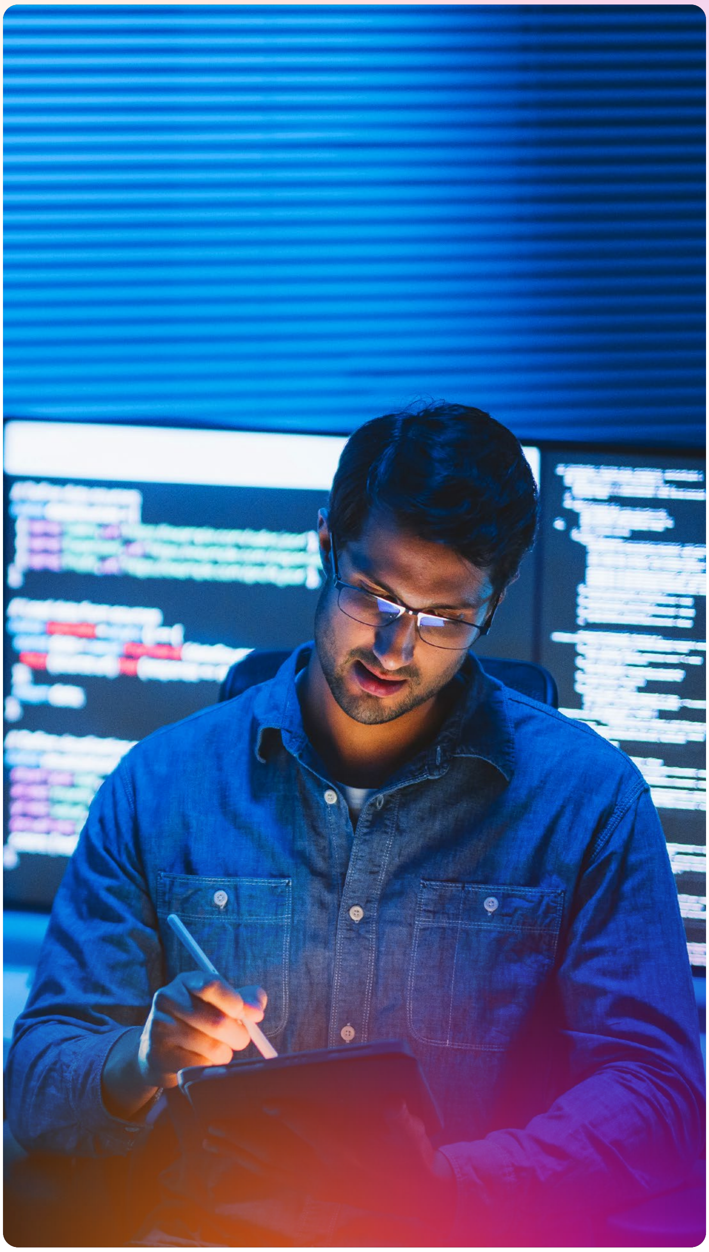

Why validate? Because it dramatically reduces false positives and wasted effort. CTEM research notes that discovery tools often produce many "low-confidence" issues based on theoretical patterns.

Prioritization narrows the list, but validation further filters it to what an attacker can do. By confirming exploitability, you avoid spending heavy resources on issues that turn out to be non-exploitable or minor. Validation also helps you predict the impact of assets as well. In Gartner's words, think of prioritization as re-ordering the to-do list, and validation as filtering that list based on attacker behavior. One produces a ranked list; the other cuts the list to the critical few.

Make sure to document the outcomes. For every top exposure, record whether it was confirmed exploitable and what the potential impact is. This builds a knowledge base and is extremely useful for communicating to executives about why certain fixes are critical. It also helps justify any acceptance of risk. If an issue is not easily exploitable or has minimal impact, you have evidence to support a decision to accept or defer it.

Step 03

Develop outcome-driven metrics and reporting

As mentioned in the introduction, one reason CTEM adoption lags is the difficulty many security leaders have in demonstrating value to stakeholders.

Classic metrics like "number of vulnerabilities patched" or compliance scores don't translate well to business risk reduction. CTEM needs bespoke, business-tuned metrics.

Done right, these metrics serve two purposes:

Internal improvement: It provides feedback on what's working or not and can highlight process bottlenecks if something is too slow.

Executive communication: They allow you to show management and the board how CTEM is making the company safer over time, justifying budget and attention. Gartner stresses that recording and reporting risk reduction impact is one of CTEM's most valuable outputs.

What to measure

Consider a mix of metrics that cover both the activity of managing exposures and the outcomes in risk terms. Some examples:

Exposure discovery & reduction metrics:

- **Number** of exposures identified in each cycle
- **Number** of exposures remediated or mitigated. Focus on closure rates
- **Mean Time** to Remediate (MTTR) exposures, from discovery to fix
- **Trends** in the above over time

Validation & impact metrics:

- **Percentage** of validated critical findings
- **Number** of Attack paths that lead to critical accounts or IP

Program maturity metrics:

- **Coverage** - percentage of enterprise IT covered by CTEM assessments
- **Stakeholder engagement** -number of business units actively participating in CTEM process

Establish a dashboard

For executives/CISO/board: A quarterly report might highlight key risk indicators, top exposures resolved, and any major residual risks. Use visuals – trending charts, heat maps – to show progress.

For operational teams: A more detailed dashboard to track open exposures by owner, aging of issues, etc., to manage the workflow. This will be bespoke to every business.

Outcome-driven metrics are an essential component of CTEM

From the start of your program, decide what success looks like and how you'll measure it. Baseline your starting point and then track improvements as you implement each step. Metrics will prove CTEM's worth.

Step 04

Mobilize all departments in remediation

Even the best prioritization and detection of exposures will fail to reduce risk if you can't quickly remediate or mitigate those exposures.

Teams need to establish strong collaboration and workflows with all the stakeholders required to fix the problems identified. This includes IT operations, application developers, cloud engineers, third-party vendors, or even non-IT units like HR or Finance.

Speak the language of the business and IT

When engaging other departments, how you communicate exposure matters. It's often counterproductive to run to another team shouting about CVEs and attack vectors. Instead, frame the issue in terms of the risk to something they care about.

For example, tell the application owner: "This vulnerability could allow an attacker to steal customer data from your application – it's like leaving a back door open. We need your team's help to close it within 2 weeks to meet our security risk appetite for customer data."

Shared tools and integrated data

Integrating CTEM findings into the everyday tools of IT and business teams is crucial. If your DevOps team uses Jira, ensure CTEM tasks are in Jira. If IT has a configuration management database (CMDB) or asset management, link CTEM data to those assets and so forth.

Don't forget third parties either. "Wide engagement" extends outside the company when needed. If a critical exposure is in a vendor's product or a supplier's system, someone (procurement, vendor management) must lean in to ensure the vendor provides a fix or mitigation. Include such external scenarios in your mobilization plan.

Flexible in mobilization

Not every exposure can be "fixed" in the traditional sense. Some may require long-term projects or acceptance. In those cases, mobilization might mean mitigating or monitoring. The key is that every significant exposure is owned and has a defined path forward – whether that's remediation, mitigation, or formally accepted risk with compensating controls.

Nothing should fall into a black hole. By embedding CTEM in this way, you move to a unified front beyond the security team against threats. This cultural shift may take time, but the benefits are immense: faster closure of exposures, fewer surprises, and a shared sense of responsibility for protecting the business.

Step 05

Iterate continuously: simultaneous scopes & ongoing posture validation

The final step of implementing a CTEM program is to make sure it's continuous.

Unlike a one-off project or annual audit, running multiple CTEM cycles concurrently, and continuously validates your security posture even as you remediate issues. Essentially, CTEM becomes part of your organization's rhythm, continuously looking for the next exposure, even while the last one is being fixed.

Simultaneous scopes

In Step 1, we discussed defining scopes and discovery and that you can have more than one scope. Here, we emphasize that CTEM operates in cycles per scope, and these cycles are typically ongoing and overlapping.

A standard CTEM cycle includes Scoping → Discovery → Prioritization → Validation → Mobilization. When one cycle ends, the next begins, perhaps with an expanded scope or refreshed discovery. You are always either discovering new exposures or fixing known ones, ideally simultaneously. If you are running multiple scopes, these can be on different cadences. You might assess critical customer-facing systems continuously and internal systems on a monthly cycle, for instance.

Ongoing posture validation

While in Step 2 we discussed validation in terms of checking specific exposures, here posture validation refers to a broader practice which includes regularly testing your overall security posture through adversarial simulations even when you think things are "patched and secure."

Essentially, this concept is continuous red teaming or continuous purple teaming. It ensures that, even if you've remediated all known exposures, your defenses and processes are ready for unknown threats or any attack that slips through. No matter how well you follow Steps 1-4, the threat landscape can change in an instant, and you might have unknown exposures.

Continuous posture validation helps reveal those unknown unknowns. By simulating attacks, you might find that you have no high CVEs open. Even so, an attacker could exploit a combination of medium issues to achieve a critical breach; something a vulnerability scan alone wouldn't flag.

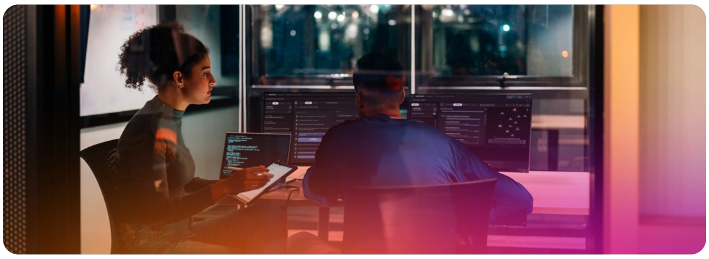

Methods for continuous validation include

- **Continuous** attack surface discovery
- **Scheduled** red/purple team exercises
- **Automated** breach and attack simulation tools running frequently to test common attack paths in the network
- **Intentionally** stress certain controls under safe conditions to ensure they work

Even as you patch and fix known issues, you're asking "What did we miss? How would we fare if attacked right now?" On a rolling basis. It closes the loop of CTEM by feeding back into Scoping and Discovery for the next cycle against any newfound weaknesses.

For example, you might run a tabletop exercise that reveals your incident response team wasn't aware of a particular cloud asset, implying a blind spot in discovery – that insight would lead you to update your CTEM scope to include that asset class, and possibly add new telemetry sources to discovery. In this way, continuous validation drives continuous improvement of the CTEM process. This fosters a mindset of cyber resilience where the organization is not just trying to prevent attacks by patching holes, but also regularly checking if an attack occurred and if the team could detect and limit it.

When done well, CTEM and incident response start to blend. In fact, one outcome is often improved incident response playbooks. By examining potential attack paths, you ensure you have plans for those scenarios and by running scopes simultaneously and continuously validating, you keep the momentum and ensure no lull in security improvement.

Conclusion

Gartner's CTEM framework represents a shift from ad-hoc vulnerability management to a continuous, risk-based security practice. By following these five steps, organizations can begin realizing the benefits of CTEM:

01 **Align exposure scopes with business priorities:** Focus your efforts where they matter most.

02 **Prioritize and validate findings:** Use risk context to prioritize, then test the highest-risk issues to confirm which ones truly demand immediate action.

03 **Develop bespoke metrics:** What gets measured gets improved and funded. Create metrics that translate to business outcomes.

04 **Mobilize a cross-functional response:** Breaking down silos and fostering shared ownership of cyber risk is vital.

05 **Make it continuous:** This is not a one-time project. Run multiple scopes in parallel and keep iterating.

For organizations early in this journey, start with a pilot on a critical scope and iterate. Early success in one domain will make the case expand CTEM across the enterprise. For more mature organizations, assess where you stand on each of the five steps. Perhaps you have vulnerability prioritization down to a science but need to improve cross-team mobilization, or vice versa.

By continuously aligning with business needs, rigorously prioritizing what matters, quantitatively tracking progress, engaging broadly, and never standing still, CTEM enables you to systematically reduce your cyber risk exposure. While this guide provides the foundational steps to get started, embracing CTEM will position your organization to not only thwart today's threats but also to adapt quickly to tomorrow's, fostering a security program that is always a step ahead to reclaim control of your security.

Prioritize on true cyber risk with Darktrace

We hope this checklist provides a practical framework to help you evaluate and create a CTEM strategy with tools that best align with your organization's needs.

Now that you have the steps in hand, we highly encourage you to familiarize yourself with Darktrace / Attack Surface Management, Darktrace / Proactive Exposure Management, and Darktrace / Incident Readiness & Recovery. Also referred to as the Cross Platform suite, all three play their part in helping teams of any size stay ahead of threats before they can affect the network.

/ Attack Surface Management ↗

Is that most external and outward facing point in looking at risk. New CVE information arises and feeds into your workflow, immediately showing which assets are potentially affected via the Newsroom feature. Alongside a rating system to easily identify at-risk assets, teams can also validate potential vulnerabilities.

/ Proactive Exposure Management ↗

Connects the outward facing information from Darktrace / Attack Surface Management with the internal connections and users of your organization. Both technology and people are evaluated, massively improving the ability to see what areas have the biggest potential business-relevant impact. Paired with realistic phishing engagements, it's incredibly easy to mobilize teams and validate findings.

/ Incident Readiness & Recovery ↗

Is an extension of Darktrace / NETWORK, residing in the primary Threat Visualizer UI. Teams are able to augment red teaming and tabletop exercises via simulations of real-world incidents superimposed onto an organization's network. For both simulations and real incidents, teams are supported by dynamic, AI-generated playbooks which help guide post-incident actions clearly. When it comes to continuous iteration and testing what works, Darktrace / Incident Readiness & Recovery is a clear addition.

■ **About Darktrace**

Darktrace is a global leader in AI cybersecurity that keeps organizations ahead of the changing threat landscape every day. Founded in 2013 in Cambridge, UK, Darktrace provides the essential cybersecurity platform to protect organizations from unknown threats using AI that learns from each business in real-time. Darktrace's platform and services are supported by 2,700+ employees who protect nearly 10,000 customers globally. To learn more, visit www.darktrace.com.

North America: +1 (415) 229 9100 Europe: +44 (0) 1223 394 100 Asia-Pacific: +65 6804 5010 Latin America: +55 11 4949 7696

darktrace.com | info@darktrace.com