

SEARCHLIGHT.
CYBER

ASIM

in the age of CTEM

MICHAEL GIANARAKIS

MJ KAUFMANN

MICHAEL GIANARAKIS

Michael Gianarakis has worked in the security industry for over a decade building and managing offensive security teams across the Asia Pacific and Japan. In 2018 he co-founded the Attack Surface Management company Assetnote, which was acquired by Searchlight Cyber in 2025. Michael is now SVP of ASM at Searchlight. Michael has presented his security research around the world including at DEF CON, Black Hat Asia, BSides, Las Vegas, Hack in the Box, AusCert, Thotcon, 44Con and OWASP.

MJ KAUFMANN

MJ Kaufmann is founder and principal consultant of Write Alchemist. With over two decades of IT and security expertise, her experience ranges from trailblazing enterprise-level projects at the University of Florida to shaping the next generation of IT professionals as an associate professor to freelancing and ghostwriting for cybersecurity and technology companies. Although MJ's business specializes in working with start-ups and helping them grow, she's also done work for global tech leaders like UST, Saviynt, BitDefender, Cisco, and Snyk on various projects.

SEARCHLIGHT.
CYBER 

ASIM

in the age of CTEM

MICHAEL GIANARAKIS

MJ KAUFMANN

CONTENTS

PART I: REDEFINING ASM	7
ASM's Unfulfilled Promise	8
The Dissolved Perimeter	10
Third-Party Risk Explosion	12
PART II: BUILDING A CTEM PROGRAM	15
Section 1: ASM Is the Foundation, CTEM Is the Process	16
Section 2: CTEM Stage-by-Stage — Powered by ASM	18
1. Scoping – Defining What to Evaluate	18
2. Discovery – Identifying Real Exposures	19
3. Prioritization – Knowing What Matters Most	19
4. Validation – Determining If an Exploit Is Feasible	20
5. Mobilization – Assigning Remediation to the Right Teams	20
Section 3: Four Critical Failure Modes When ASM Is Missing	22
1. Misaligned Scope	22
2. False Positives and Alert Fatigue	22
3. Missed Exposures	23
4. Inefficient Remediation	23
Section 4: Capabilities of a CTEM-Ready ASM Program	24
1. Continuous Asset Discovery	24
2. Exploit-Based Validation	25
3. Deep Asset Enrichment	25
4. Ownership and Context Mapping	26
Section 5: CTEM Maturity Starts with ASM Discipline	28

PART III: INTELLIGENCE-DRIVEN ASM —	31
TURNING SIGNALS INTO STRATEGY	
Section 1: Moving Beyond Generic Threat Intelligence	32
The Problem with Most Threat Intel	32
ASM + Threat Intelligence = Contextual Detection	33
Key Capabilities of Intelligence-Driven ASM	34
Section 2: The Vulnerability Intelligence Gap	36
Time-to-Exploit vs. Time-to-Remediate	36
ASM’s Role in Closing the Gap	36
Building a Vulnerability Intelligence Pipeline	37
Real-World Example: MOVEit Transfer RCE	38
Section 3: Operationalizing Research-Driven Security	40
Offensive Security for Defensive Programs	40
Capabilities to Build or Buy	40
Real-World Applications	41
Section 4: Creating Feedback Loops Across CTEM	42
Detection-Informed Discovery	42
Threat-Informed Prioritization	42
Enrichment Loops	43
Section 5: Scaling Intelligence with Automation and Integration	44
Integrate Intelligence into the CTEM Workflow	44
Use API and Custom Checks to Automate Intelligence Consumption	45
Tie Back to Business Risk	46
Section 6: The CTEM Intelligence Maturity Model	48
CONCLUSION	51



PART



Redefining ASM

ASM'S UNFULFILLED PROMISE













When Attack Surface Management (ASM) first entered the security vocabulary, it promised to provide a continuous, external view of organizational assets. On paper, it would function as an always-on inventory of everything attackers could see and potentially exploit. In practice, however, the category often underdelivered.

The core issue wasn't the idea itself, but how it was implemented. Early ASM tools focused heavily on asset discovery but stopped short of verifying whether those assets actually introduced meaningful risk. The result was an overwhelming flood of information with little prioritization. Security teams were handed lists of domains, IP addresses, and ports—a digital census with no risk hierarchy. Valuable time was lost sorting signal from noise.

This disconnect between theory and execution turned ASM into a passive cataloging exercise rather than an active security function. The assumption was that discovery alone would drive better outcomes. But without verifying which assets were vulnerable or exploitable, and without contextualizing those assets within the business, ASM became just another data source to ignore.

The missing piece was exposure verification. Knowing something is exposed is different from knowing it is exploitable. That distinction is everything when security resources are limited and attacker dwell time is measured in hours. Exposure verification should have been central from the start: not just identifying potential risks, but proving whether and how they could be used against you.

Knowing something is exposed is different from knowing it is exploitable. That distinction is everything when security resources are limited and attacker dwell time is measured in hours.

THE PROMISE OF ASM	THE REALITY OF ASM
 Continuous Asset Discovery	 Discovery only
 Exposure Verification	 Static Asset Lists
 Contextual Risk Prioritization	 No Context or Risk Scoring
 Asset Ownership Resolution	 Surface-Level Insight
 Integrated Threat Intel	 No Verification
 Automated Testing	
 Security Feedback Loop	

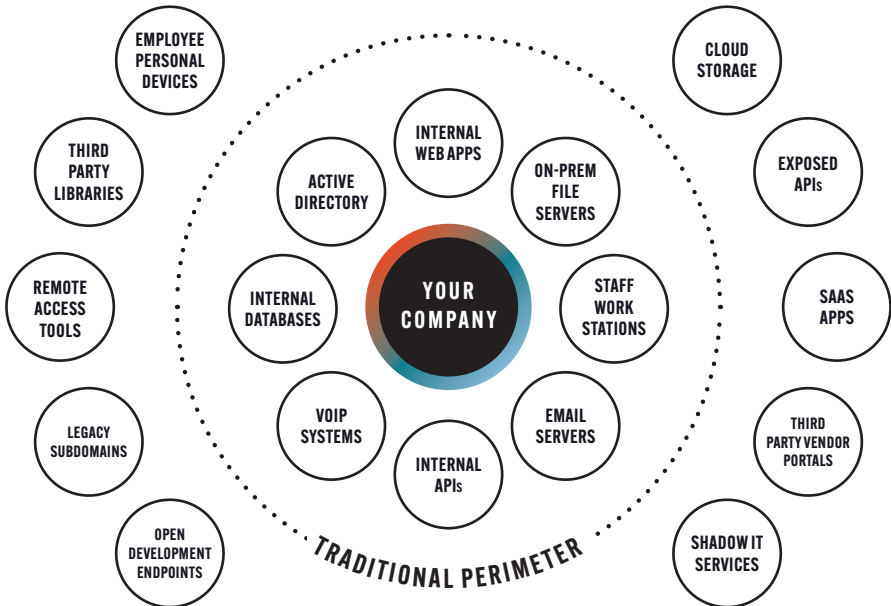
This is where the modern ASM model pivots. Rather than relying on static scans or surface-level analysis, forward-looking ASM approaches integrate continuous testing and embedded vulnerability research. It's not just about discovering assets; it's about operationalizing security research to understand how those assets could be targeted. This creates a feedback-driven system where findings are validated, prioritized, and directly actionable.

The integration of techniques from continuous penetration testing and offensive research closes the vulnerability intelligence gap that exists between vendor disclosure and attacker exploitation. By identifying exposures before they become headlines, ASM becomes a living, responsive system—one that proactively helps organizations reduce their risk profile, not just document it.

THE DISSOLVED PERIMETER

The historical model of cybersecurity was built around the concept of the perimeter: a clearly defined boundary between trusted internal networks and the untrusted external world. But cloud adoption, remote work, and rapid development practices have erased that boundary.

Modern attack surfaces are sprawling, dynamic, and borderless. Organizations no longer control all the infrastructure their data passes through. A single business unit might rely on a mix of SaaS tools, APIs, third-party platforms, and microservices deployed across multiple clouds. This ecosystem changes constantly—with assets appearing and disappearing as development cycles spin up new features or retire old ones.



The cloud transformation amplified this shift. Infrastructure-as-Code (IaC) and automation mean new assets can be deployed and exposed in seconds, often without security oversight. Temporary workloads, ephemeral containers, and serverless functions increase complexity and decrease visibility. Security teams can't protect what they don't know exists.

Shadow IT compounds the challenge. When departments onboard new SaaS tools or spin up cloud resources without informing central IT or security, those assets fall outside of governance and monitoring. These blind spots introduce unmanaged risk, especially when they handle sensitive data or integrate with core systems.

APIs, once an afterthought, have become a major attack vector. As systems become more interconnected, APIs expose functionality and data that can be targeted if not secured. Their surface is often dynamic, versioned, and inconsistently documented—making them difficult to track and harder to defend. An exposed development endpoint or misconfigured API gateway can be just as dangerous as an open RDP port.

Finally, microservices and containerization bring their own visibility hurdles. Each service may live independently, scale dynamically, and communicate across multiple environments. Legacy scanning tools struggle to keep up with these architectures, often missing transient services or failing to map dependencies between components. The result is a fragmented view that attackers can exploit and defenders may not even see.

The dissolved perimeter isn't a theoretical shift—it's the operating reality. And it demands a shift in mindset. ASM must map what the internet sees, not what internal teams think they own. That means real-time discovery, frequent verification, and a model that adapts as quickly as the infrastructure it monitors.

THIRD-PARTY RISK EXPLOSION

As organizations lean into agility, collaboration, and outsourcing, their exposure expands in directions they can't always control. Third-party risk has become a dominant security concern—not just because vendors can be compromised, but because organizations often have limited visibility into how their systems are connected.

Supply chain attacks have moved from rare occurrences to standard operating procedures for sophisticated threat actors. Vulnerabilities in vendor systems can quickly propagate to customers. Integration points become attack paths. And the larger your partner ecosystem, the harder it is to monitor and secure it all.

Vendor APIs and plugins often create direct pathways into sensitive environments. A misconfigured or compromised integration can provide attackers with authenticated access without ever breaching the “front door.” These integration points are particularly dangerous because they are often trusted by default and monitored inconsistently.

Third-party JavaScript libraries present another problem. Every external script embedded on a webpage increases the attack surface. Malicious updates or supply chain compromises in these libraries can affect thousands of organizations at once. The risk isn't theoretical—we've seen attackers compromise CDN-hosted libraries to distribute malware from legitimate websites.

Third-party risk has become a dominant security concern—not just because vendors can be compromised, but because organizations often have limited visibility into how their systems are connected.

Open source dependencies offer scale and speed to developers, but at the cost of increased security complexity. Many applications rely on outdated or unvetted packages that may contain known vulnerabilities or insecure logic. Dependency chains grow long and tangled, making it difficult to trace risk back to its origin. An attacker doesn't need to breach your infrastructure if they can breach your software supply chain.

What makes third-party risk especially dangerous is its diffuse ownership. Security teams may not have insight into what vendors or libraries are in use. Even if they do, they may lack the access or authority to enforce remediations. When a zero-day hits a popular dependency or SaaS platform, organizations often find themselves racing to understand if they're even affected.

Managing this complexity requires expanding the boundaries of ASM. It's no longer enough to scan your known IP blocks or audit internal systems. External attack surface management must extend into supplier, vendor, and integration ecosystems. That includes monitoring for newly exposed third-party systems, detecting abnormal changes in vendor-facing assets, and correlating asset exposure with known risks in the supply chain.

Visibility must reach beyond your organization to include everything connected to it. The dissolved perimeter isn't just about cloud and containers. It's also about trust boundaries and dependencies. And ASM's evolution depends on illuminating the connections most likely to be ignored.



PART



Building a CTEM Program

SECTION 1

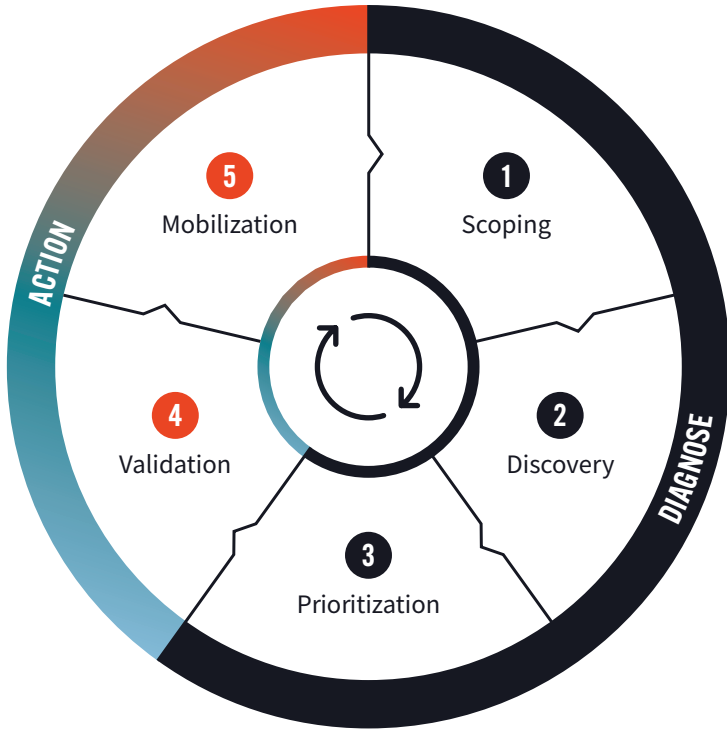
ASM IS THE FOUNDATION, CTEM IS THE PROCESS

Continuous Threat Exposure Management (CTEM) offers a strategic framework for identifying and mitigating risk, but it can only be as effective as the data it relies on. Without comprehensive, real-time visibility into the external attack surface, CTEM efforts are incomplete from the outset. This limitation isn't just about CTEM. As organizations move toward broader external cyber risk management strategies, blind spots in the attack surface can undermine any attempt to quantify or mitigate risk across digital assets.

Security teams often rush to implement simulations, validations, and automation workflows without first establishing an accurate, continuous view of their external exposures. This results in wasted cycles, false confidence, and missed vulnerabilities. ASM must come first. It is the map that defines what CTEM can explore. And when that map is incomplete or outdated, every subsequent action becomes less reliable.

Continuous Threat Exposure Management (CTEM) offers a strategic framework for identifying and mitigating risk, but it can only be as effective as the data it relies on.

GARTNER'S FIVE STAGES OF CTEM



SECTION 2

CTEM STAGE-BY-STAGE – POWERED BY ASM

Gartner defines CTEM as comprising five stages: Scoping, Discovery, Prioritization, Validation, and Mobilization. Each of these depends on a strong ASM foundation to function effectively.

Gartner defines CTEM as comprising five stages: Scoping, Discovery, Prioritization, Validation, and Mobilization.

1

SCOPING

DEFINING WHAT TO EVALUATE

Scoping determines which systems and assets are in focus for exposure management activities. If this is based solely on known IP ranges or legacy inventories, it leaves critical gaps.

ASM enables high-fidelity, real-time asset discovery across cloud providers, SaaS environments, APIs, shadow IT, and subsidiary infrastructures. This ensures that scoping decisions are comprehensive and informed.

Example: A security team launches a CTEM initiative scoped to their AWS account and internal IP blocks. However, they overlook a third-party SaaS integration used by the finance team that stores sensitive customer data. That SaaS platform later becomes a breach vector because it was never evaluated.

2

DISCOVERY

IDENTIFYING REAL EXPOSURES

Many organizations still rely on point-in-time, IP-centric scans that can't keep up with the pace of cloud deployments or infrastructure-as-code. As a result, critical misconfigurations or vulnerabilities often go undetected.

ASM's continuous, hourly discovery mechanisms ensure that new and modified assets are captured as soon as they are exposed. This includes transient cloud resources, newly opened ports, or ephemeral APIs that traditional scanners miss.

Key differentiator: ASM enables detection of exposures as they happen—not hours later, not after a breach, and not buried in logs.

3

PRIORITIZATION

KNOWING WHAT MATTERS MOST

CTEM is designed to reduce noise and help teams focus on what's truly exploitable. But prioritization is only meaningful if it's based on validated risk, not theoretical vulnerability scores.

ASM helps by providing:

- **Exploit-Based Verification:** Proof-of-concept exploits that demonstrate real-world feasibility.
- **Business Context:** Asset tagging and ownership mapping that align exposures with their potential impact.

Without this context, CTEM workflows often collapse under the weight of triage and alert fatigue.

4 VALIDATION

DETERMINING IF AN EXPLOIT IS FEASIBLE

Red teaming, breach and attack simulation (BAS), and automated pen testing tools are powerful, but only if they target relevant, exposed infrastructure. ASM ensures that validation efforts are based on real exposures, not assumptions.

It also provides:

- Up-to-date asset context for chaining simulations
- Support for identifying lateral movement paths

You can't validate what you haven't mapped.

5 MOBILIZATION

ASSIGNING REMEDIATION TO THE RIGHT TEAMS

Even the most accurate findings are useless if no one knows who owns the asset. Without ownership metadata, security alerts sit idle while exposure windows stay open.

ASM contributes by:

- Tagging assets by business unit or responsible team
- Routing notifications to the correct people
- Enabling automated ticket creation and follow-up

Example: A verified exposure is discovered on a legacy domain. It gets flagged in the dashboard but ignored for weeks because no team claims ownership. Meanwhile, attackers exploit it to pivot deeper into the network.



CTEM is designed to reduce noise and help teams focus on what's truly exploitable.

But prioritization is only meaningful if it's based on validated risk, not theoretical vulnerability scores.

SECTION 3

FOUR CRITICAL FAILURE MODES WHEN ASM IS MISSING

When organizations move forward with CTEM initiatives without a strong ASM foundation, they encounter predictable and often costly failure modes. These issues don't just affect efficiency—they directly impact the organization's ability to detect, validate, and remediate real threats.



MISALIGNED SCOPE

CTEM relies on accurate scoping to determine what systems and assets are included in the exposure management process. Without ASM, those scopes are often incomplete or outdated. Critical elements like SaaS applications, third-party APIs, or forgotten subdomains are left out entirely.

This results in blind spots—entire portions of the attack surface that go unmanaged. When threat actors exploit these gaps, defenders are caught off guard, reacting to breaches they didn't know were possible.



FALSE POSITIVES AND ALERT FATIGUE

CTEM tools that ingest raw, unvalidated data from legacy systems end up creating noise instead of clarity. When exploit validation is missing, every minor misconfiguration or informational alert can appear urgent.

Over time, security teams lose confidence in the tools themselves. Analysts grow slower to respond, and important alerts blend in with background noise. This fatigue delays real responses and introduces unacceptable risk windows.



MISSED EXPOSURES

Shadow IT, unmanaged third-party services, and ephemeral cloud assets are all examples of attack surface components that change constantly. Without continuous ASM, these assets are often unknown, untracked, and unprotected.

Missed exposures mean missed threats. And when attackers find those exposures before defenders do, the consequences are immediate: data loss, access escalation, and brand damage.



INEFFICIENT REMEDIATION

Even if a CTEM system identifies a legitimate vulnerability, that insight is meaningless without context. If no one knows who owns the asset, how critical it is, or which business unit it supports, remediation stalls.

Ownership confusion leads to delays. Vulnerabilities sit in dashboards without resolution. Meanwhile, attacker dwell time increases, and exposure windows widen.

Even if a CTEM system identifies a legitimate vulnerability, that insight is meaningless without context. If no one knows who owns the asset, how critical it is, or which business unit it supports, remediation stalls.

SECTION 4

CAPABILITIES OF A CTEM-READY ASM PROGRAM

To make CTEM actionable, ASM must evolve from static discovery to a continuous, context-rich, validation-driven capability. A CTEM-ready ASM program delivers four critical functions:

To make CTEM actionable, ASM must evolve from static discovery to a continuous, context-rich, validation-driven capability.

1

CONTINUOUS ASSET DISCOVERY

Discovery can't happen once per quarter, week, or even day. Modern environments change by the minute, so ASM must offer:

- Agentless, external-first discovery
- Passive and active reconnaissance
- Integration with cloud providers, including wildcard domain resolution

This model ensures complete visibility into dynamic infrastructure like containers, serverless workloads, ephemeral virtual machines, and API gateways.

2 EXPLOIT-BASED VALIDATION

Real risk doesn't come from theoretical vulnerabilities—it comes from exposures that can be exploited. ASM must:

- Use proof-of-concept exploits and scripting to test feasibility
- Confirm vulnerabilities with real-world techniques drawn from security research
- Reduce false positives and boost team trust in findings

Exploit validation transforms data into decisions. It keeps security teams focused and accelerates remediation cycles.

3 DEEP ASSET ENRICHMENT

Raw IP addresses and DNS records aren't enough. ASM must add context to each asset, including:

- Technology fingerprints (client and server)
- Software versions and SSL certificate details
- API endpoints and behavioral metadata

Historical change records improve incident response, while visual indicators—like screenshots, titles, and request logs—provide security teams with rapid situational awareness.

4**OWNERSHIP AND CONTEXT MAPPING**

Remediation happens faster when responsibilities are clear. ASM should:

- Tag assets by team, business unit, location, and sensitivity
- Link exposures to internal ownership for faster triage
- Support automated ticket routing and escalation paths

This capability closes the loop between detection and action. When everyone knows who's responsible, vulnerabilities don't get lost in translation.

Together, these four pillars enable ASM to support every phase of CTEM—from initial scoping to final remediation—while continuously improving detection quality, team efficiency, and overall security outcome.

CHECKLIST FOR A CTEM-READY ASM PROGRAM

- Continuous Asset Discovery**
Hourly scanning of your infrastructure including cloud providers.
- Exploit-Based Validation**
Use of Proof-of-concept exploits and scripting to test feasibility.
- Deep Asset Enrichment**
Go beyond IP addresses and DNS records to include technology fingerprints, software versions, and API endpoints
- Ownership and Context Mapping**
Exposures linked to internal ownership to enable faster and more efficient triage

SECTION 5

CTEM MATURITY STARTS WITH ASM DISCIPLINE

Building a CTEM program without first investing in ASM is like trying to build a skyscraper on sand. The structure might take shape, but it won't hold.

ASM isn't a checklist item or a one-off inventory effort—it's a continuous system that informs and improves every step of exposure management. When ASM is treated as a foundational discipline rather than a reactive function, organizations gain clarity, speed, and resilience.

Mature ASM programs deliver measurable advantages. These same characteristics are foundational for external cyber risk management, where organizations are expected to measure and reduce risk exposure across assets they own, operate, or depend on.

FOUR METRICS FOR MEASURING CTEM PERFORMANCE



Mean Time to Remediation (MTTR):

How long it takes security teams to remediate confirmed exposures in your infrastructure.



Mean Time of Exposure (MTE):

The duration between an exposure appearing in the attack surface and security teams detecting it.



Remediation Velocity:

How quickly teams resolve high-priority vulnerabilities compared to industry benchmarks.



Coverage of Asset Discovery:

The percentage of internet-facing assets on your network being actively monitored.

These advantages are:

- **Faster Mean Time to Remediation (MTTR):**
Because exposures are identified, verified, and routed with context, they're resolved more quickly.
- **Fewer False Positives:**
Exploit-based validation cuts through noise, giving teams confidence in their data.
- **Better Executive Reporting:**
With clear asset ownership, tagging, and exposure histories, reporting becomes accurate, consistent, and tied to business priorities.

CTEM maturity isn't just about adopting the latest tools or frameworks. It begins with visibility—the ability to see what's exposed, understand how it could be exploited, and know who can fix it.

You can't secure what you can't see. ASM provides that visibility. CTEM gives you velocity. But one doesn't work without the other.

Building a CTEM program without first investing in ASM is like trying to build a skyscraper on sand. The structure might take shape, but it won't hold.



PART



Intelligence-Driven ASM
Turning Signals into Strategy

SECTION 1

CTEM is more than a process — it's a living system that improves over time. That improvement depends on embedding intelligence into every step: threat intel, research, exploit verification, and contextual enrichment. This part explores how to operationalize those capabilities.

MOVING BEYOND GENERIC THREAT INTELLIGENCE

The Problem with Most Threat Intel

Many organizations rely on mass-distributed intelligence feeds—lists of IPs, URLs, file hashes, and domains that have been flagged somewhere as suspicious. These may include generic indicators of compromise (IOCs), blocklists, or IP reputation scores. While useful in bulk filtering, they offer very little contextual relevance to a specific environment.

These feeds tell you what has been bad for someone, somewhere. But they don't tell you whether it's bad for you. That disconnect leads to alert fatigue, wasted effort, and missed prioritization. Without knowing whether a flagged IP is interacting with your environment—or what it's interacting with—security teams are left guessing.

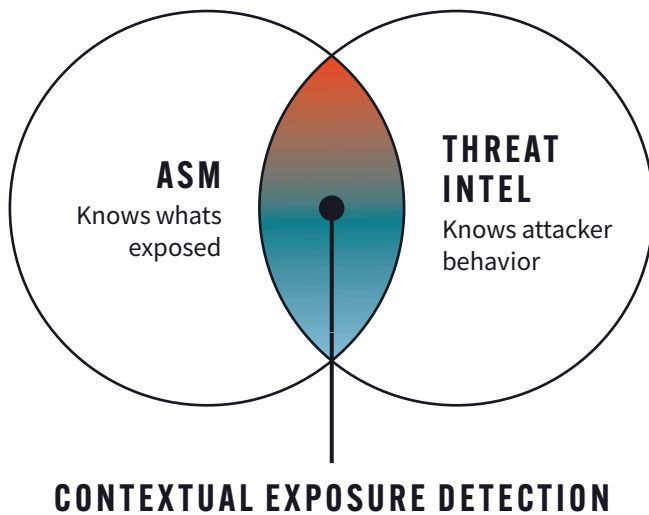
ASM + Threat Intelligence = Contextual Detection

This is where ASM becomes the missing link. By continuously mapping what's truly exposed across your external attack surface, ASM grounds threat intelligence in operational context. It turns abstract signals into concrete risk assessments.

An IOC becomes actionable only when it intersects with something real. For example, a flagged IP from a threat feed is relevant if—and only if—it's communicating with a known-exposed admin panel in your infrastructure. ASM tells you whether that panel exists, where it's hosted, and what business unit owns it.

The combination of ASM and threat intelligence produces context-aware detection:

- ASM knows which assets you're exposing.
- Threat intel identifies attacker tactics.
- Together, they show which attackers are likely to succeed—and where.



Key Capabilities of Intelligence-Driven ASM

Integrating threat intelligence with ASM enables a shift from hypothetical security posture to real-time adversary alignment. Several key capabilities support this approach:

- **Correlate Exposures with Threat Campaigns:** Match newly discovered vulnerabilities to known threat actor techniques and campaigns. For example, if a ransomware group is known to target a specific file transfer application, ASM can check for exposures of that software across your environment.
- **Identify Tech Stack Patterns That Match Known APT Behavior:** Threat groups often focus on specific technologies. ASM helps security teams spot where those stacks are deployed and whether they're exposed to the internet. This turns threat intelligence into a lens, not just a list.
- **Enable Threat Hunting Rooted in Your Actual Surface:** Hunting is more effective when it starts with known exposures. ASM makes it possible to filter for vulnerabilities that match an attacker's tactics, techniques, and procedures (TTPs), focusing investigation on where attackers are most likely to go.

By transforming generic signals into asset-specific threat indicators, intelligence-driven ASM bridges the gap between knowledge and action. It ensures your CTEM program isn't just aware of global threats—but tuned to the ones that matter most to you.

“

CTEM is more than a process — it’s a living system that improves over time. That improvement depends on embedding intelligence into every step: threat intel, research, exploit verification, and contextual enrichment.

SECTION 2

THE VULNERABILITY INTELLIGENCE GAP

Time-to-Exploit vs. Time-to-Remediate

When a zero-day vulnerability becomes public, attackers often move faster than defenders. Weaponization can happen within hours of disclosure, long before most organizations are able to patch. CTEM platforms that depend on vendor advisories or CVSS scores to initiate response processes are already behind.

The delay between discovery and defensive action creates a critical exposure window. Every hour matters.



ASM's Role in Closing the Gap

Continuous ASM solutions with embedded research capabilities can detect vulnerable software before any public disclosure occurs. By monitoring changes in open-source projects, analyzing patch diffs, and developing proof-of-concept exploits, integrated research teams are able to identify exploitable patterns proactively.

This kind of embedded intelligence helps customers act earlier—sometimes weeks or months before vendors release official guidance.

Building a Vulnerability Intelligence Pipeline

A mature ASM program supports a structured pipeline for ingesting and operationalizing vulnerability intelligence. That pipeline includes:

Source Signals:

- Proprietary research, including zero-day identification
- Open-source intelligence (OSINT)
- Exploit marketplaces and underground forums
- Patch diffing and reverse engineering

Ingestion:

- Feed these signals directly into the ASM system
- Match against asset fingerprints and exposed services
- Enrich findings with severity ratings, exploit maturity, and known attacker interest

Automation:

- When a new CVE is linked to your tech stack, trigger an immediate scan across the attack surface
- Initiate validation routines, issue alerts, or launch pre-configured mitigation workflows

This kind of automation tightens the loop between discovery and action, turning research into immediate protection. As more organizations adopt external cyber risk management frameworks, this feedback loop becomes essential—not just for finding issues, but for quantifying and controlling risk across distributed infrastructure.

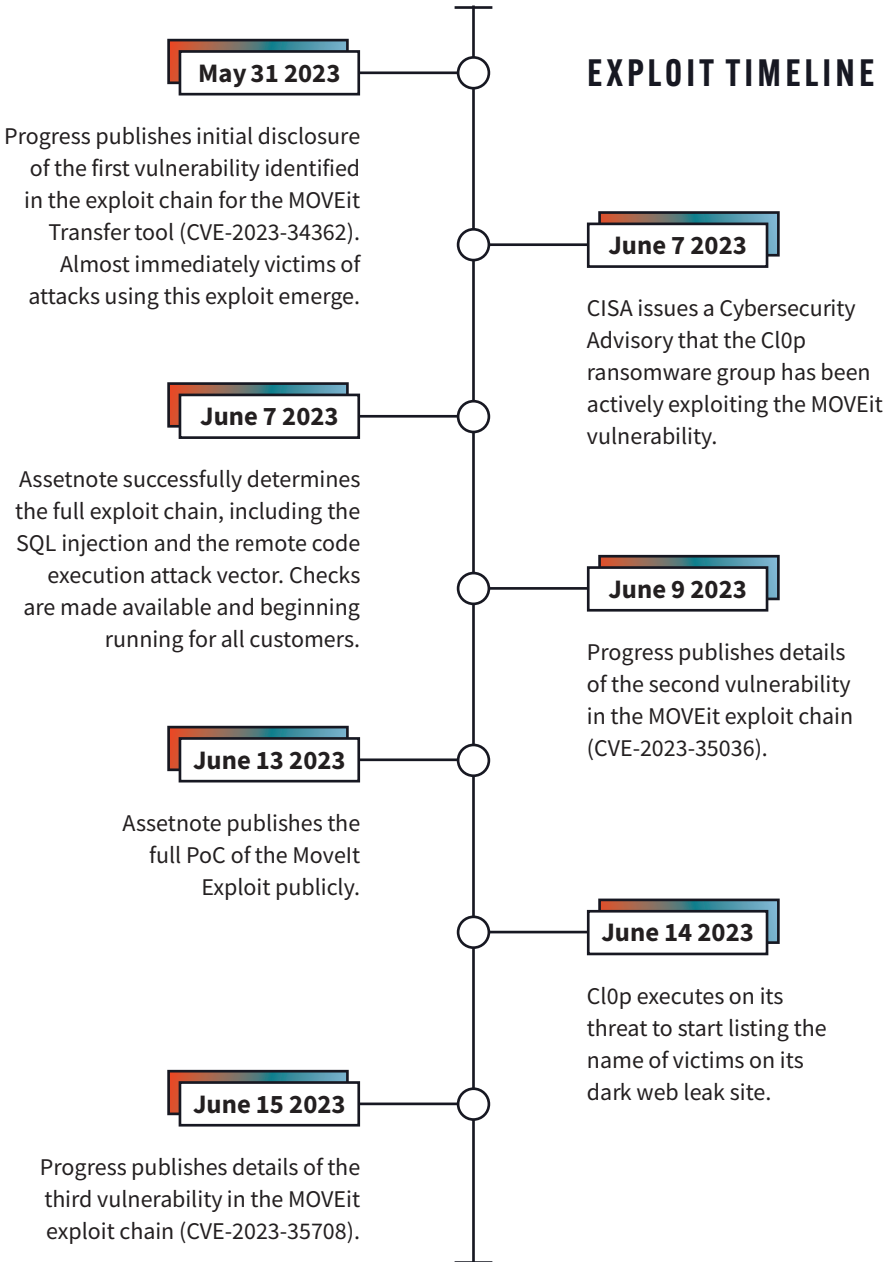
Real-World Example: MOVEit Transfer RCE

A practical example of this model in action came during the MOVEit Transfer remote code execution vulnerability exploited by ransomware groups in mid-2023. While most organizations first became aware of the vulnerability after the CVE was published and attacks began, customers using platforms with embedded vulnerability research had detections already in place.

According to Assetnote's technical breakdown , their researchers had identified the vulnerability through independent analysis before the vendor issued public disclosure. That research was operationalized into detection logic and deployed within their ASM platform. Customers were able to identify exposure and act before exploitation campaigns reached scale.

This kind of response isn't theoretical—it's the direct result of combining exploit research, ASM discovery, and intelligence automation. And it's exactly the model organizations need to close the gap between vulnerability emergence and defensive readiness.

EXPLOIT TIMELINE



SECTION 3

OPERATIONALIZING RESEARCH-DRIVEN SECURITY

Offensive Security for Defensive Programs

Threat actors operate with precision, chaining vulnerabilities and moving laterally through environments to reach their targets. Defensive teams must adopt the same mindset. CTEM is not just a compliance exercise; it's a continuous simulation of what real attackers would do. To be effective, exposure validation must mimic the logic of offensive campaigns.

Treating CTEM like a hybrid red/blue team operation ensures that findings are meaningful. Exploit development, chaining techniques, and lateral movement insights all inform stronger defenses when applied to detection and validation. Each exposure should be tested not just for existence, but for feasibility. This adversarial approach helps organizations understand how a vulnerability could be used, what it could connect to, and what damage it might enable.

Capabilities to Build or Buy

To make this model real, security teams need both talent and tooling. Whether built internally or sourced from vendors, certain capabilities are essential:

Security Research Expertise:

- Conduct static and dynamic analysis to dissect code, understand behavior, and identify attack vectors.
- Reverse-engineer patches to uncover silently fixed vulnerabilities or logic errors.
- Go beyond CVEs to detect post-authentication flaws and custom application weaknesses that scanners miss.

CTEM is not just a compliance exercise; it's a continuous simulation of what real attackers would do. To be effective, exposure validation must mimic the logic of offensive campaigns.

Platform Support:

- Create and manage custom vulnerability signatures that reflect real-world attacker logic
- Deploy PoCs to auto-validate exposures as they are discovered
- Detect anomalies without relying on signatures by monitoring for unusual asset behavior, such as unexpected service changes or open ports on normally quiet assets

Together, these capabilities make research-driven ASM feasible at scale. They allow CTEM teams to shift from passive detection to predictive analysis.

Real-World Applications

Operationalizing research isn't about building a lab for its own sake—it's about producing outcomes. Some practical use cases include:

- **Exploit PoCs for Validation:** Before triggering remediation processes, use working exploit code to confirm the risk is real. This reduces false positives and speeds up MTTR by focusing on what actually needs fixing.
- **Simulating Chaining Scenarios:** Many serious breaches result from chaining together multiple small exposures. Internal threat labs can simulate these chains to understand how a seemingly low-severity issue could lead to privilege escalation or lateral movement.
- **Detecting Regressions Over Time:** By tracking exploitable logic changes across asset updates, organizations can detect when a previously fixed vulnerability reappears. This capability is especially important in CI/CD environments, where frequent updates increase the chance of accidental reintroduction.

SECTION 4

CREATING FEEDBACK LOOPS ACROSS CTEM

CTEM programs should get better with time. That means not just collecting intelligence, but feeding it back into the system to improve detection, prioritization, and remediation processes.

Detection-Informed Discovery

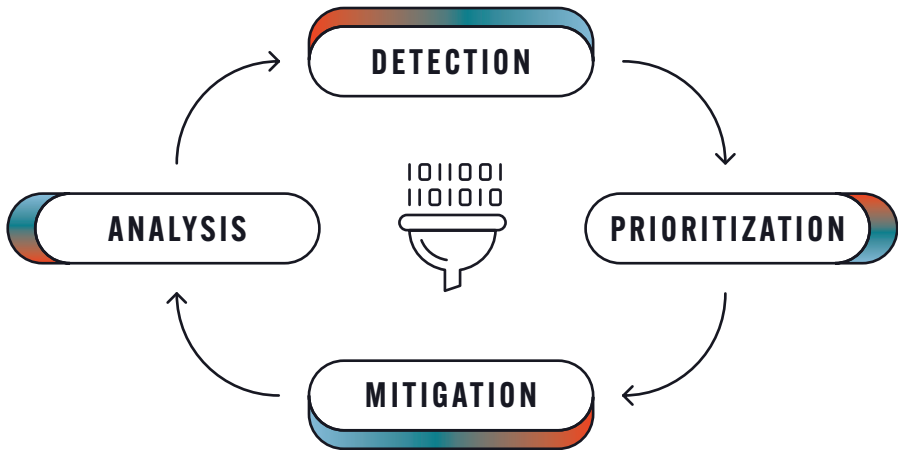
Every validated exposure is a chance to learn. Security teams should use confirmed findings to develop reusable patterns and signatures. Feeding those patterns into the ASM platform helps detect similar exposures across other business units, environments, or subsidiaries.

For example, if a specific SaaS integration exposes an internal dashboard, that exposure pattern can be modeled and searched for elsewhere.

Threat-Informed Prioritization

Not all vulnerabilities are equal. Prioritization should reflect what attackers are actually doing—not just what scanners flag. Intelligence about in-the-wild exploitation should raise the urgency level of corresponding exposures. On the other hand, vulnerabilities that remain theoretical, or have no working exploit, may be deprioritized to preserve analyst bandwidth.

An intelligence-driven feedback loop ensures that every discovery, validation, and resolution strengthens the system.



Enrichment Loops

Closed exposures provide valuable context. By analyzing how and why an incident occurred, teams can refine asset tagging, detection logic, and validation techniques. For example:

- A previously unknown asset that became a breach vector might reveal a gap in tagging or team attribution.
- An exploit chain might highlight a missing dependency in the prioritization model.

By systematically integrating insights from past exposures, CTEM becomes smarter. Detection becomes sharper. Response becomes faster.

Outcome: An intelligence-driven feedback loop ensures that every discovery, validation, and resolution strengthens the system. Instead of reacting to exposures in isolation, organizations build a CTEM program that improves with every cycle.

SECTION 5

SCALING INTELLIGENCE WITH AUTOMATION AND INTEGRATION

CTEM maturity depends not only on the quality of intelligence gathered, but also on how well that intelligence is integrated and automated across security workflows. Scaling research-driven insights requires a seamless connection between data, detection, and decision-making.

Integrate Intelligence into the CTEM Workflow

To move from manual threat tracking to proactive defense, CTEM programs should directly integrate with Threat Intelligence Platforms (TIPs) and external enrichment sources. This integration enables:

- Auto-prioritization of new findings based on active threat signals, such as indicators flagged for current exploitation campaigns
- Triggered validation workflows, where intelligence cues launch automated PoC execution, exploit module checks, or asset-specific scans

For example, if a new TTP associated with a known threat actor matches the behavior of a vulnerable API endpoint, the ASM platform can auto-run a verification routine and elevate the issue for immediate attention.

This level of integration removes guesswork and lets threat data drive action within seconds, not hours.

Use API and Custom Checks to Automate Intelligence Consumption

Advanced ASM platforms offer flexible APIs and custom scanning capabilities to ingest external data sources and act on them in real time. This includes:

- **Custom Signatures for Zero-Days:** As new vulnerabilities are discovered by internal or vendor research teams, organizations can quickly create custom checks and scan their attack surface for exposures
- **Secret Detection in Code Repositories:** Automatically monitor GitHub, GitLab, and similar services for leaked secrets or credentials using specialized detectors
- **Real-Time Feed Ingestion:** Match threat feeds like abuse.ch or the CISA Known Exploited Vulnerabilities (KEV) catalog against the current attack surface to flag any matching assets

Automation in these scenarios dramatically increases responsiveness. Instead of waiting for tickets or blog posts, detection and validation happen the moment intelligence is received.

Tie Back to Business Risk

All this automation and enrichment must tie back to what matters most—protecting the systems that power the business. This means:

- **Mapping intelligence to business-critical applications:** If an exploit targets a specific CMS or framework, ASM should quickly identify whether it's present on high-value assets
- **Feeding executive dashboards with context-rich insights:** Rather than technical summaries, deliver metrics like: “Three of our externally exposed apps are running software targeted by active ransomware campaigns.”

This linkage between threat intelligence and business impact ensures the right people understand not just what is vulnerable, but why it matters. For security leaders advancing toward external cyber risk management, this connection between operational exposure and strategic business risk is the foundation for setting priorities and allocating resources effectively.

The result is a CTEM program that doesn't just scale in size—it scales in value. By integrating intelligence directly into workflows, automating detection, and tying risk to business context, organizations build a system that continuously adapts to both attackers and objectives.



CTEM maturity depends not only on the quality of intelligence gathered, but also on how well that intelligence is integrated and automated across security workflows.







SECTION 6

THE CTEM INTELLIGENCE MATURITY MODEL

Building a resilient CTEM program is not a one-step transformation. It's a maturity journey. Organizations move through stages as they enhance how intelligence is collected, contextualized, and applied. This progression begins with basic IOC matching and evolves into adaptive systems capable of auto-validating exposures, prioritizing based on real-world threats, and informing executive strategy in real time. The chart below outlines a practical framework for assessing and advancing CTEM intelligence maturity.

CTEM only delivers on its promise when it's grounded in continuous discovery and driven by contextual intelligence. ASM provides the visibility. Intelligence sharpens the focus. Together, they form the foundation of a modern exposure management strategy—one that learns from every signal and strengthens with every response.

CTEM only delivers on its promise when it's grounded in continuous discovery and driven by contextual intelligence. ASM provides the visibility. Intelligence sharpens the focus.

LEVEL	CAPABILITY	DESCRIPTION
	Fully Adaptive, Intelligence-Oriented CTEM System	Auto-prioritization, auto-validation, executive visibility, real-time learning
	Feedback-Informed Improvement	Every detection and resolution makes ASM/CTEM smarter
	Research-Driven Detection	Acting on pre-disclosure findings, building custom checks
	Exploit Validation	Incorporating PoCs to confirm exploitability
	Asset-Centric Intel	Mapping intel to known assets
	Generic Feeds	Consuming vendor threat reports, basic IOC matching

A dark background with a light-colored grid of dots. Two orange lines with circular endpoints are positioned on the left side. One line starts from the left edge and extends diagonally upwards to a circle. The other line starts from the left edge and extends horizontally to a circle.

Conclusion

CONCLUSION

Continuous Threat Exposure Management is a relatively new concept in the wider context of cybersecurity but it stands on the shoulders of established practices. In many ways, it can be seen as a continuation of the Attack Surface Management movement that began over half a decade ago.

Certainly, it has the potential to fulfill the promise of what ASM should have been: establishing a continuous, external view of the risk inherent in an organization's infrastructure. CTEM is also a useful mechanism to help organizations adapt their security practices to modern realities, both in terms of the technological makeup of their attack surface and the tactics that are used by hackers today.

However, ensuring the potential is realized this time round will require change. Companies will need to increase the frequency of their asset discovery, which means recognizing that weekly or monthly scans really don't fit the bill of "continuous asset discovery". They will also need to move to exploit-based validation of vulnerabilities, to ensure the signal of a true threat isn't drowned out by the noise of a thousand false alarms.

Most importantly of all, successful implementation of CTEM requires an acceptance that it is a process and not a tool. Of course, it would be easier if we could just tell you to plug in the best CTEM solution and press "play", but that entirely misses the point. The power of CTEM is inherent in the fact that it is a framework, not a prescribed technology, which means it has the flexibility to adapt as technology and threats evolve.

Adopting this shift in mindset will also allow security teams to overcome one of the other greatest challenges in cybersecurity: measurability.

The clear metrics of CTEM - remediation time, exposure time, and asset coverage - provide practitioners with something tangible they can bring to the board in language they can understand.

The real benefits, of course, can't be measured. They are the vulnerabilities that aren't exploited. The attacks that don't happen. The data and intellectual property that isn't lost. All prevented through a more proactive approach to cybersecurity.

The power of CTEM is inherent in the fact that it is a framework, not a prescribed technology, which means it has the flexibility to adapt as technology and threats evolve

SEARCHLIGHT.
CYBER 

VISIT WWW.SLCYBER.IO/ASM TO FIND
OUT MORE OR **BOOK A DEMO NOW.**

UK HEADQUARTERS

Suite 63, Pure Offices,
1 Port Way, Port Solent,
Portsmouth PO6 4TY
United Kingdom

US HEADQUARTERS

44 Merrimac Street,
Newburyport,
MA 01950
United States