

EXECUTIVE WHITEPAPER | 2026 EDITION

BUILDING YOUR CTEM PROGRAM STEP BY STEP

A practical guide for CISOs, SOC leaders, and security architects on operationalizing Continuous Threat Exposure Management from first step to full programme maturity.

SCOPE

Unified Vulnerability Management · Threat Informed Defense · Unified Exposure Validation · Risk Quantification

AUDIENCE

CISOs · SOC Leaders · Security Architects · MSSPs

FRAMEWORK

CTEM · MITRE ATT&CK® · NIST CSF 2.0 · NIST 800-53 · NIST RMF · NIST 800-30 · NIST 800-39

01

Executive Summary

The case for Continuous Threat Exposure Management in 2026

In 2026, the cybersecurity landscape has undergone a fundamental shift. The attack surface grows by an estimated **23% annually**, AI agents now discover zero-day vulnerabilities daily, and adversaries execute multi-stage attack paths using Living Off The Land (LOTL) techniques that evade traditional defenses. Yet most organizations continue to rely on static vulnerability scoring, periodic penetration tests, and siloed tools that generate alerts without context.

The result is **"reactive exhaustion"**, security teams paralyzed by an alert swarm, spending 80% of their time triaging incidents instead of preventing the attacks that actually threaten the business. The critical question every CISO must answer "Can a real adversary, using known techniques, execute an attack in our environment and will we detect it?" remains unanswered by most security stacks.

The CTEM framework provides the answer: a continuous, business-aligned cycle that scopes, discovers, prioritises, validates, and mobilises remediation transforming security from a series of disconnected technical tasks into a strategic, preemptive defense.

This whitepaper is a practical guide for organizations ready to move beyond the scan-and-patch cycle. It maps the complete CTEM journey across four integrated modules Unified Vulnerability Management (UVM), Threat Informed Defense (TID), Unified Exposure Validation (UEV), and Unified Risk Management & Quantification (URMQ) and shows how each module builds on the last to create a self-reinforcing, measurable risk reduction engine.

95%

of "Critical" vulns are not reachable by an attacker

40%

of EDR/SIEM/firewall controls fail silently due to misconfiguration

60%

reduction in MTTR achievable through AI-driven agentic workflows

23%

annual growth rate of the enterprise attack surface in 2026

WHAT THIS WHITEPAPER COVERS

Chapter 1 — The Problem

Why traditional VM, BAS, and point-in-time pentesting can no longer keep pace with the modern threat landscape.

Chapter 5 — Exposure Validation

Proving your controls actually work with agentic purple teaming, BAS, and Detection as Code.

Chapter 2 — The CTEM Framework

The five-step CTEM cycle and how KnightGuard operationalizes each stage with AI-native & Risk-centric workflows.

Chapter 6 — Risk Quantification

Translating technical security posture into quantifiable business risk for CISO and board reporting.

Chapter 3 — Starting with UVM

How to establish a unified 360-degree exposure intelligence layer as the foundation of your CTEM program.

Chapter 7 — The Feedback Loop

How the four modules form a self-reinforcing cycle of continuous risk reduction.

Chapter 4 — Threat Informed Defense

Operationalizing adversary intelligence moving from theoretical risk to TTP-aligned defensive coverage.

Chapter 8 — Building the Business Case

ROI, TCO, and the metrics that matter for justifying CTEM investment at the executive level.

02

The Problem: Three Systemic Gaps

Why traditional security approaches fail against modern adversaries

Traditional vulnerability management was built for a static era on-premises servers, scheduled scans, annual penetration tests. It was never designed for today's hyper-connected, cloud-native enterprise. Security teams are now paralyzed by an **"Alert Swarm"** of thousands of high-severity CVEs with no business context, treating a printer firmware flaw with the same urgency as a critical database misconfiguration. The gap between discovering a flaw and its actual exploitation has shrunk dramatically yet the manual effort to validate and remediate has not. The result is alert fatigue: critical exposures remain unaddressed, buried under mountains of non-exploitable noise.

Modern adversaries do not exploit isolated vulnerabilities they execute **Multi-Stage Attack Paths**, leveraging real-world Tactics, Techniques, and Procedures (TTPs) across an organization's infrastructure, often using legitimate tools (Living Off The Land). This creates three fundamental gaps that drive the need for CTEM:

GAP 01 — THE CONTEXT GAP

Too many vulnerabilities, too little relevance

Standard scanners identify millions of flaws, but less than 5% of "Critical" vulnerabilities are actually reachable by an attacker. Without validation, teams waste 95% of effort on non-exploitable noise exhausting resources with zero risk reduction.

GAP 02 — THE DETECTION BLIND SPOT

Unproven MITRE ATT&CK coverage

40% of EDR, SIEM, and firewall controls fail silently due to misconfiguration or drift. Organizations assume their security stack is defending as promised but without continuous validation against known TTPs, that assumption is dangerous and unfounded.

GAP 03 — THE EXECUTION GAP

Disconnected offense and defense

Traditional BAS tools operate in a silo, separate from vulnerability data and remediation workflows. Security teams might prove a breach is possible but lack the automated pipeline to instantly harden defenses leaving a critical "execution gap" that adversaries exploit.

THE IMPACT OF THESE GAPS

IMPACT AREA	CONSEQUENCE
Signal-to-Noise Ratio	High operational overhead processing non-exploitable alerts
Mean Time to Remediate	Increased MTTR as teams lack context to prioritise correctly
Security Investment ROI	Underutilization of SIEM, EDR, and threat intel platforms
Business Risk Quantification	Inability to translate technical exposure into board-level risk
Adversary Advantage	Attackers exploit validated paths while defenders chase ghosts

THE FUNDAMENTAL MISALIGNMENT

Security teams measure exposure based on theoretical weaknesses CVSS scores, patch counts, number of critical alerts. Adversaries exploit validated, contextual attack paths across an organisation's real infrastructure, using known TTPs that existing controls frequently miss. CTEM closes this gap by shifting the measurement to what actually matters: exploitability, business impact, and validated control effectiveness.

THE STRATEGIC SHIFT CTEM ENABLES

Reactive → Preemptive

CVSS Scores → Business Impact

Siloed Alerts → Unified Risk View

Assumed Security → Validated Defence


Patch Counts → Measurable Resilience

03

The CTEM Framework: Five Steps to Continuous Resilience

How KnightGuard operationalizes the CTEM cycle with AI-native & Risk-centric workflows

Continuous Threat Exposure Management (CTEM) is a framework that moves security from point-in-time assessments to a continuous, business-aligned cycle. It scopes and prioritises the most relevant threats and assets, enabling organizations to dynamically reassess remediation efforts as environments, technologies, and threat vectors change. KnightGuard operationalizes all five stages through a unified platform with embedded Agentic AI giving CISOs a single pane of glass across the entire exposure lifecycle.



STEP-BY-STEP: WHAT HAPPENS AT EACH STAGE

1

Scoping — Defining the Battleground

Configure ASM to map Crown Jewels and internet-exposed perimeter

- Map digital footprint: domains, subdomains, shadow IT, exposed services, etc.
- Define target risk profile and map to the current state to identify gaps
- Quantify risk based on potential financial impact to the business
- Eliminate noise by focusing exclusively on business-relevant risk

2

Discovery — Uncovering the Invisible

Enable multi-cloud and dark web connectors

- Identify misconfigurations across AWS, GCP, Azure, GitHub, Kubernetes, etc.
- Dark web monitoring: leaked credentials, brand mentions, hacker chatter
- Map current controls to MITRE ATT&CK and NIST 800-53
- Aggregate findings from all VM scanners into a unified source of truth

3

Prioritization — Cutting Through the Noise

Activate the AI-Enabled Ingestion Pipeline

- AI normalizes BYO threat intel and maps it to internal discovery data
- Move beyond CVSS: integrate KEV, EPSS, exploit probability, asset value
- Identify kill chain paths which vulnerabilities lead to crown jewels *
- Choke Point Analysis: fix one path, block ten possible attacks
- Adversary aligned defence prioritisation

* In road map

4

Validation — Proving the Defence

Launch Unified Purple Teaming and Agentic Hunting

- 1,500+ pre-built emulation campaigns test real-world breach scenarios
- Red Team AI Agent executes attacks; Blue Team Agent audits SIEM response
- AI-enabled detection as Code: generate, test, and deploy SIEM rules in minutes
- Prove which controls work replace assumption with evidence

5

Mobilization — Driving the Outcome

Integrate with ITSM and generate verified outcomes

- Bi-directional sync with Jira, ServiceNow, Slack, and etc.
- Every prioritised risk gets a ticket, a remediation playbook, and a tracked outcome
- Auto-map remediations to compliance frameworks.
- CISO dashboards show exactly how actions reduce the Threat Exposure Score

THE KNIGHTGUARD AGENTIC LAYER

Unlike static platforms, KnightGuard deploys a Mesh of AI Agents CTI Analyst, SOC Analyst, Red & Blue Team, and ITOps that possess Reasoning and Action (ReAct) capabilities. They query your SIEM, update tickets, generate detections, and write remediation playbooks autonomously, operating as a 24/7 force multiplier for your security team.

04

Module 1: Unified Vulnerability Management

Start here to build the 360-degree exposure intelligence foundation

The journey begins with **Unified Vulnerability Management (UVM)** the strategic foundation that unifies your entire exposure landscape into a single source of truth. UVM does not rip-and-replace your existing investments. It sits as a unification layer on top of Tenable, Qualys, and cloud-native scanners augmented by AI Agents to centralize and operationalize all security intelligence.

THE 360-DEGREE EXPOSURE INTELLIGENCE LAYER



KEY UVM CAPABILITIES

- Attack Surface Management (Internal & External)**
Continuous outside-in discovery of your digital footprint domains, subdomains, shadow IT, exposed services, SSL certificates, etc.
- Cloud Misconfigurations** NEW
Real-time agentless auditing across AWS, GCP, Azure, GitHub, and Kubernetes. Drift detection identifies deviations from Standard Benchmarks. Auto-tags against SOC2, HIPAA, GDPR.
- Multi-Scanner Normalization**
Ingests all leading VM scanners (commercial and open source). De-duplication merges redundant findings, reducing data noise by up to 60% one clean record per unique vulnerability.
- AI-Driven Vulnerability Prioritization**
Beyond CVSS: integrates KEV, EPSS, dark web intel, exploit probability, threat actor activity, and asset value into a dynamic multi-factor risk score contextualized to your specific environment.
- Vulnerability Remediation AI Agent**
Transforms "what is wrong" into "how to fix it" providing exact PowerShell scripts, GPO settings, and tool-specific instructions. Contextual guidance avoids breaking business dependencies.
- Bidirectional ITSM Integration**
Real-time sync with Jira, ServiceNow, Slack, etc. When a ticket is updated, KnightGuard reflects the status instantly shifting from "alerting" to "executing."

UVM VS. TRADITIONAL VM: THE DIFFERENTIATORS

CAPABILITY	SUPPORTED FRAMEWORKS	KNIGHTGUARD UVM
Intel Operationalization	Reliant on public feeds (CISA KEV) or static vendor data. Slow to adapt.	AI pipeline ingests unstructured BYO-Intel (PDFs, blogs, advisories) into actionable logic in minutes.
Remediation Effort	High-level "Apply Patch" recommendations. Technical how-to left to IT team.	ITOps AI Agent generates tool-specific, executable playbooks tailored to your exact OS and security stack.
Risk Scoring	Linear CVSS + manual asset tagging. Fails to reflect dynamic threat landscape.	Proprietary algorithm dynamically weights scores using real-world simulation outcomes and validated control status.
Tool Agnosticism	Creates vendor lock-in, pushing users toward proprietary scanner ecosystems.	Fully agnostic mesh layer consolidates ROI of existing scanners without rip-and-replace.

05

Module 2: Threat Informed Defense

Operationalizing adversary intelligence to defend against what matters

With your exposure foundation in place, the next step is to inject **Adversary Intelligence** into every security decision. Threat Informed Defense (TID) shifts the focus from defending against "everything" to defending against "what matters" the specific TTPs known to target your industry, geography, and technology stack. By leveraging MITRE ATT&CK®, GIR, and BYO threat intelligence, KnightGuard TID ensures your security posture is dictated by real-world adversary behavior, not guesswork.

WHY THREAT-INFORMED DEFENSE?

The majority of security breaches do not stem from zero-day vulnerabilities they stem from **Gaps in Threat-Informed Defence**. Attackers rely on known TTPs that leave detectable patterns. Focusing on identifying and disrupting these behaviors enables the shift from reactive detection to proactive resilience.

MITRE ATT&CK Alignment

Real-time heat map of coverage against 200+ adversary techniques. Identifies exactly which techniques your controls currently detect, which are partially covered, and where critical gaps exist.

GIR-Based Intelligence Prioritization

General Intelligence Requirements (GIR) framework categorizes threat intelligence that transforms raw cyber underground data into prioritised, actionable intelligence answering the "who, what, when, where, why, and how" of threats targeting your organisation.

BYO Threat Intelligence

Ingest and instantly apply your existing commercial threat intel subscriptions. The CTI Analyst AI Agent normalizes unstructured advisories, blogs, and proprietary reports into structured, actionable ATT&CK mappings in minutes.

Controls Prioritization

Multi-framework mapping across NIST 800-53, MITRE Mitigation, and MITRE D3FEND automatically identifying which controls to implement first for maximum risk reduction across prioritized TTPs.

One-Click Detection Operationalization

Identify the most relevant detection rules for prioritized TTPs and push them directly to your SIEM/EDR with a single click in Sigma, FQL, KQL, Splunk, or any query language.

TID 12-POINT COVERAGE MODEL

INTELLIGENCE

- Threat Intel & Operationalization
- Threat Prioritization
- MITRE ATT&CK Matrix

CONTROL MANAGEMENT

- Controls Prioritization
- NIST CSF Mapping
- Controls Remediation

DETECTION

- Threat Monitoring
- Detection Analytics
- AI-Native Hunt Analytics

AGENTIC AI

- AI Agent: CTI Analyst
- AI Agent: IT Ops
- AI-enabled detection as code

TID AGENTIC AI USE CASES

CTI Analyst AI Agent — Automated TTP Mapping

When a new ransomware strain is identified in the wild, the CTI Agent automatically parses the threat report, identifies relevant MITRE ATT&CK Techniques, and highlights specific gaps in current security controls — converting a news headline into a defensive action within minutes, not days.

EXPECTED OUTCOMES FROM TID IMPLEMENTATION

60%

reduction in MTTD & MTTR through AI-enabled playbooks



Cross-functional synergy between GRC, SecOps, and ITOps

06

Module 3: Unified Exposure Validation

From vulnerability visibility to adversary reality proving your controls actually work

Knowing where you're exposed is not enough. The critical question every security team must answer is: "Can a real adversary, using known attack techniques, successfully execute an attack in our environment and will we detect it?" KnightGuard's Unified Exposure Validation (UEV) module answers this definitively by combining 1,500+ real-world emulation campaigns with AI-driven Detection Engineering creating a high-speed feedback loop that transforms offensive validation from a simple "test" into an automated "upgrade" for your defensive stack.

THE PURPLE TEAM FEEDBACK LOOP

CONSUME	VALIDATE	ENGINEER	ORCHESTRATE
Ingest Threat Data	Red Team Tests	Detection as Code	Push & Patch
High-risk ATT&CK and techniques from the TID module	AI Agents generate exploits for each stage	AI Agents generate context specific detections	Rule to SIEM, ticket to ITSM for patching

KEY UEV CAPABILITIES

Agentic Purple Teaming

The Red agent helps generate exploits for specific scenarios while the Blue agent simultaneously generates detection analytics to validate control effectiveness replacing weeks of manual work with minutes of automated precision.

Detection as a Code (DaaC) NEW

Treat your security logic like software. Manage detections life cycle within the platform. Write the rule once; automatically translate it for Splunk, Sentinel, QRadar, or Elastic simultaneously.

1,500+ Ready-to-Deploy Campaigns

Extensive library of adversary playbooks mapped to the latest threat intelligence, updated weekly. Covers lateral movement, data exfiltration, credential theft, and LOTL scenarios without disrupting production environments.

Multi-SIEM Unification NEW

Eliminate the "language barrier." Normalize detection logic across Splunk, Sentinel, QRadar, and Elastic deploying a single rule across all SIEMs simultaneously.

UEV VS. TRADITIONAL BAS: THE DIFFERENTIATORS

CAPABILITY	STANDARD BAS	KNIGHTGUARD UEV
Validation Method	Static path analysis (mathematical)	Active emulation: 1,500+ real-world campaigns
SIEM Support	Usually proprietary or single-connector	Multi-SIEM unification: one rule, many platforms
Logic Management	Manual rule creation in console	Detection as Code: scalable, versioned, automated
AI Integration	Chatbots for "summarization"	Agentic AI: autonomous Red/Blue agents that execute tasks
Threat Hunting	Reactive, IOC-based	Proactive: intelligence-led, and TTP-based

THE SECURITY STACK PARADOX - SOLVED

Organizations have invested heavily in EDR, SIEM, and firewalls yet 40% of these controls fail silently due to misconfiguration or drift. UEV is the mechanism to continuously "ping" these controls, proving they defend as promised. Rather than assuming your SIEM will catch an exploit, KnightGuard runs the emulation, confirms whether the alert fires, and if not generates the exact detection logic to fix it instantly.

ADVANCED THREAT HUNTING

KnightGuard's pre-built hunting queries leverage the same intelligence used in emulation campaigns, generating hunt logic in KQL, AQL, SPL, or any query format automatically enabling proactive search for silent threats that have already evaded perimeter controls.

UEV BUSINESS BENEFITS

↓ MTTR

AI workflows automate step-by-step remediation playbooks

↑ ROI

Identify which controls work and which are redundant

07

Module 4: Unified Risk Management & Quantification

Translating technical security posture into defensible business risk

The final module in the KnightGuard CTEM journey converts everything discovered, prioritized, and validated into **Quantifiable Business Risk** the language boards, regulators, and executives understand. Rather than reporting "number of patches applied," the URMQ module gives CISOs a defensible, weighted Risk Score grounded in validated control effectiveness, real-world exploitability, and organizational context.

CYBER RISK QUANTIFICATION METHODOLOGY

KnightGuard's risk quantification is built on NIST standards NIST RMF, NIST 800-30, NIST 800-39, and NIST 800-53 providing a recognized, defensible framework for communicating risk to regulators, auditors, and boards.

Ready-to-Quantify Risk Scenarios NEW

Pre-built quantification models for 9+ risk scenarios like: BEC, Ransomware, Data Loss, and Insider Threat scenarios. No manual modelling required instantly calculate potential financial impact based on your security posture and threat profile.

Organization Context-Based Quantification NEW

Moves beyond generic industry benchmarks. KnightGuard ingests your specific technology stack, business processes, and regulatory obligations to produce risk scores that reflect your actual exposure, not theoretical averages.

Risk Gap Analysis NEW

Identifies the delta between your current security posture and your target risk profile (NIST, DORA, ISO, Ransomware-specific). Provides a prioritized roadmap of controls to implement to close the gap.

Use-Case Specific Risk Profiles

Dedicated profiles for DORA compliance, OT/ICS environments, ransomware scenarios, and sector-specific regulatory frameworks (SEBI-CSCRF, SEC, GDPR) with NIST 800-53 as the underlying framework ensuring relevance across regulated industries.

Automated Compliance Coverage

Maps remediation outcomes automatically to AWS, GCP, and Azure compliance posture. Every validated fix is recorded against the relevant framework control, building an audit-ready evidence trail.

CISO & BOARD REPORTING

REPLACING "NUMBER OF PATCHES" WITH "RISK REDUCTION"

KnightGuard transforms security from a technical cost center into a risk management function. CISO dashboards show exactly how today's actions reduced the organization's Threat Exposure Score providing the data needed to justify budget, meet regulatory requirements (DORA, SEC, SEBI-CSCRF), and report progress to the board in terms of Resilience rather than raw vulnerability counts.

FRAMEWORK ALIGNMENT

FRAMEWORK CATEGORY	SUPPORTED FRAMEWORKS	KNIGHTGUARD APPLICATION
Security Governance	NIST 800-53 & NIST-CSF 2.0	Automatically maps identified exposures to NIST control families, prioritizing fixes that satisfy regulatory audits.
Intelligence Planning	GIR (General Intelligence Requirements)	Uses GIR framework to categorize threat intelligence.
Adversary Behavior	MITRE ATT&CK®	Operationalize, defences aligned to adversary behaviours across ITOps and SecOps
Defensive Mapping	MITRE D3FEND™ MITRE MITIGATION	Suggests specific technical countermeasures (File Analysis, Decoy Environment) to neutralize identified threats.
Risk Quantification	NIST RMF, NIST 800-30, NIST 800-39	Prebuilt quantification modules for 9+ risk scenarios like BEC, Ransomware, Data Loss, and Insider Threat scenarios.

08

The Integrated Feedback Loop: How All Four Modules Work Together

A self-reinforcing cycle of continuous, measurable risk reduction

The true power of KnightGuard's CTEM approach emerges when all four modules operate as an integrated system. Each module amplifies the other module feeds, TID-prioritized TTPs drive UEV emulation campaigns, UEV validation results harden detections and update remediation priorities, and URMQ quantifies the aggregate risk reduction for board reporting. This creates a **Self-Healing, Continuously Improving Security Posture** that evolves as fast as the adversaries.

THE STRATEGIC SHIFT CTEM ENABLES

MODULE 1

Unified Vulnerability Management

- 360° exposure visibility
- Asset intelligence
- Multi-scanner normalization
- Choke Point Analysis

FEEDS RISK CONTEXT → TID

MODULE 2

Threat Informed Defense

- TTP alignment
- MITRE ATT&CK mapping
- Controls prioritization
- Detection analytics

FEEDS TTP PRIORITY → UEV

MODULE 3

Unified Exposure Validation

- 1,500+ emulation campaigns
- Agentic purple teaming
- Detection as Code
- Multi-SIEM deployment

FEEDS VALIDATED GAPS → URMQ

MODULE 4

Unified Risk Management

- Risk quantification
- Board reporting
- Compliance mapping
- Risk gap analysis

FEEDS PRIORITIES → BACK TO UVM

"The result is a resilient, self-healing system that evolves as fast as the adversaries not just a series of scans."

THE DAY IN THE LIFE OF A CISO — BEFORE VS. AFTER CTEM

BEFORE

Morning : Scoping

Team logs into multiple consoles. 1,000s of "Critical" alerts with no way to know which are exploitable.

AFTER

Morning : Scoping

One dashboard. KnightGuard correlates internet-exposed assets with cloud misconfigs and dark web leaks automatically.

BEFORE

Midday : Prioritization

New CVE drops. SecOps spends hours checking the stack. IT Ops and SecOps argue about what to patch first.

AFTER

Midday : Prioritization

AI CTI Agent maps CVE to org's tech stack flags only few assets that are actually exploitable.

BEFORE

Afternoon : Validation

Team assumes SIEM will catch exploits. Detection rules untested for 6 months. Manual red-teaming too expensive.

AFTER

Afternoon : Validation

One click: emulation confirms SIEM failed to alert. Specific KQL/SPL logic provided and deployed instantly.

BEFORE

End of Day : Reporting

CISO spends 2 hours massaging Excel data to explain to the board why the risk score hasn't moved.

AFTER

End of Day : Reporting

Platform auto-generates reports showing exactly how today's actions reduced the Threat Exposure Score by X%.

THE CTEM MATURITY JOURNEY

1

Foundation (Month 1-3) : [Deploy UVM](#)

Establish the unified exposure intelligence layer. Connect VM scanners, cloud providers, and dark web monitoring. Begin de-duplication and initial choke point analysis. Achieve single source of truth.

2

Intelligence (Month 3-6) : [Add TID](#)

Inject threat intel operationalization. Map controls to MITRE ATT&CK. Deploy BYO intel feeds. Begin one-click detection deployment to SIEM. Align remediation to NIST CSF 2.0.

3

Validation (Month 6-9) : [Add UEV](#)

Launch agentic purple teaming. Run first emulation campaigns against prioritized TTPs. Implement Detection as Code. Achieve multi-SIEM unification. Prove control effectiveness with evidence.

4

Quantification (Month 9-12) : [Add URMQ](#)

Implement cyber risk quantification. Build board-ready dashboards. Map all activities to various compliance requirements. Achieve defensible risk reduction reporting.

★

Full CTEM Maturity (12 months+)

Continuous, self-reinforcing loop. Security posture evolves in real-time with the threat landscape. Risk reduction is measurable, defensible, and reported at board level in business terms.

09

The Business Case & Conclusion

Measurable outcomes, ROI, and the path forward

Reduction in MTTR

60%

AI-driven workflows automate step-by-step remediation playbooks across SecOps, ITOps, and GRC teams

Noise Eliminated

95%

Less than 5% of "Critical" vulns are reachable. UVM focuses teams on the exposures that actually matter

Controls Failure Rate

40%

UEV continuously validates SIEM and EDR controls exposing the silent failures costing you security ROI

Blocks 10 paths

1 FIX

Choke Point Analysis identifies the single strategic fix with the broadest blast radius against attack paths

Analyst Hours Reclaimed

100S

Hours previously wasted on redundant triage redirected to strategic analysis and threat hunting

Zero to Visibility

<60 MIN

Cloud-native SaaS no infrastructure, no complex setup. Agentic onboarding delivers visibility in under an hour

THE SUBSCRIPTION MODULE APPROACH: START MODESTLY, SCALE CONFIDENTLY

KnightGuard is designed for the reality that every organization's CTEM journey starts from a different point. Customers enjoy complete flexibility to begin with just one or two modules and expand as the program matures no big-bang deployment, no lengthy professional services engagement.

CONCLUSION: BUILDING YOUR CTEM JOURNEY

In an increasingly complex adversary-driven threat landscape, building a CTEM program anchored in Threat-Informed Defense is no longer optional it is foundational to cyber resilience. By continuously identifying, prioritising, and validating exposures in the context of real-world threats and business impact, organizations can focus remediation efforts where they matter most.

A mature CTEM approach unifies security, IT, and risk teams; operationalizes threat intelligence; and adapts to change shifting security from reactive activity to continuous, measurable risk reduction. KnightGuard is designed precisely for this shift: a vendor-agnostic, AI-native platform that operationalizes the full CTEM framework through four integrated modules, each delivering standalone value while amplifying the others.

The fundamental question every CISO must answer :

"Are we actually protected against this specific threat right now?"
KnightGuard is a unified platform that merges offensive validation with defensive engineering, proves the attack path, tests the detection, writes the fix, and tracks the outcome providing a definitive answer, continuously.

ABOUT GAMBIT CYBER

Gambit Cyber B.V. is a Netherlands-headquartered cybersecurity company pioneering AI-driven, risk-centric Continuous Threat Exposure Management (CTEM). Through its AI-native and risk-centric Preemptive threat exposure management platform, **KnightGuard**, Gambit Cyber helps enterprises move from reactive, alert-driven security to continuous, validated, and threat-informed defence that reduces real-world cyber risk.

The company's core leadership team brings over 100 years of combined experience in cybersecurity across global technology organizations, having built and scaled large-scale security programs in complex, regulated environments. Gambit Cyber was recently recognized as one of **Europe's Top Cybersecurity Startups** by the **European Cyber Security Organization (ECSO)**.

Gambit Cyber is backed by a strong investor base including **Expeditions**, **Bitdefender**, **Voyager Ventures**, and seasoned angel investors. Their confidence in our vision not only fuels our growth but also brings invaluable expertise and strategic support to accelerate our journey.

OUR INVESTORS



READY TO START YOUR CTEM JOURNEY?



Gambit Cyber partners with CISOs, SOC leaders, and MSSPs to design, operationalize, and scale CTEM programs using a structured, step-by-step approach powered by KnightGuard.



Visit Our Website
www.gambitcyber.org



Chat to Sales
sales@gambitcyber.org



Chat to support
support@gambitcyber.org



Our Locations
The Netherlands | India