

WHITEPAPER · CTEM

CTEM, from the Gartner framework to an operational program

A technical guide to Continuous Threat Exposure Management. It covers the five phases, six public cases analyzed critically, common pitfalls, adoption data, success metrics and consolidated references. Vendor-independent material.

CATEGORY

**Continuous Threat Exposure
Management**

AUDIENCE

CISO, Risk, VulnMgmt, Architecture

EDITION

2026

READING

17 pages

CLASSIFICATION

Public

PRIMARY SOURCE

Gartner, CTEM framework

EXECUTIVE SUMMARY

CTEM is a continuous program, articulated across multiple technical capabilities

A synthesis of the Gartner framework, market projections and the thesis of this whitepaper.

Whitepaper 2026
CSURFACE

PROJECTED BREACH REDUCTION

3x

organizations with CTEM are 3x less likely to suffer a breach by 2026 ^{G1}

REDUCTION IN SUCCESSFUL ATTACKS

≥ 50%

with CTEM and cross-business mobilization by 2028 ^{G2}

NON-TECHNICAL EXPOSURES BY 2028

> 50%

will stem from identity, SaaS, third-party and process weaknesses ^{G3}

What CTEM is

Continuous Threat Exposure Management (CTEM) is a continuous operating model for reducing security risk, organized into **five iterative phases**, Scoping, Discovery, Prioritization, Validation and Mobilization, defined by Gartner as the natural evolution of vulnerability management, ASM and control-validation programs.

The distinction from earlier approaches lies in the **continuous orchestration** of these technologies around business priorities, rather than CVE lists ranked by severity.

The program requires instruments able to operate at the cadence of the threat, discover exposures beyond the obvious, prioritize by real exploitability and validate remediation. Implementation combines technical capabilities already present in the organization with new capabilities needed to close coverage gaps.

How to read this document. Pages 3 and 4 establish the paradigm and an overview of the five phases. Pages 5 to 9 detail each phase of the framework. Pages 10 to 12 present public cases, common implementation pitfalls and adoption data. Pages 13 and 14 cover maturity and metrics. Page 15 maps CSURFACE capabilities to the cycle. Pages 16 and 17 present references and next steps.

The thesis of this whitepaper

CTEM is consolidating as an operational category because the acceleration of exploitation has made the point-in-time scan cycle unworkable. Implementing CTEM requires a combination of capabilities, continuous discovery, threat intel, empirical validation and integration with mobilization, orchestrated around business priority. This whitepaper presents the framework in a vendor-independent way, and at the end shows how CSURFACE can contribute.

Gartner references. ^{G1} Gartner, "Implement a Continuous Threat Exposure Management Program (CTEM)", 2026 projection. ^{G2} Gartner, 2028 projection. ^{G3} Gartner, 2028 projection.

Executive Summary

The Paradigm Shift
The Five Phases
01 · Scoping
02 · Discovery
03 · Prioritization
04 · Validation
05 · Mobilization
Public Cases
Common Pitfalls
CTEM Adoption
Program Maturity
Success Metrics
CSURFACE → CTEM
References
Next Steps

From vulnerability management to exposure management

Why the periodic-scan model inherited from the 2000s stopped protecting, and how CTEM replaces it.

The inherited model: point-in-time vulnerability management

For two decades, the vulnerability management function operated a familiar cycle. Periodic scan, CVE list, CVSS-based prioritization and a remediation queue. The model was born when the window between disclosure and exploitation was measured in weeks or months, and when the scope to protect was a well-defined set of servers and workstations.

Today this model has three structural limitations.

- **The pace of the threat has changed.** In 2024, more than 150 CVEs were exploited within five days of disclosure. The monthly scan delivers a snapshot that is already out of date.
- **The scope of what is exposed has changed.** SaaS applications, cloud environments provisioned by engineering, acquired brands and embedded third-party scripts do not appear in a CMDB or an internal scan.
- **The prioritization criterion has changed.** Static CVSS does not distinguish what is theoretically severe from what is under active exploitation. Human triage capacity is spent on alerts that will never be attacked.

When Gartner formalized CTEM in 2022 and 2023, the objective was to describe an operating model that **connects discovery, prioritization and validation in a continuous cycle driven by business risk**, rather than by technical inventory or abstract severity.

The turning point. "Vulnerability management" describes a *function*. CTEM describes a *program*. The function seeks to eliminate known vulnerabilities. The program reduces business exposure, including weaknesses that never appear as a CVE: identity, configuration, process, digital supply chain.

The cost of the old paradigm. In large organizations, more than 250,000 open vulnerabilities coexist with a remediation capacity of only 10%. The exhaustive-queue paradigm is mathematically unworkable. CTEM replaces "everything" with "what matters, validated".

What changes in practice

From inventory to exposure. An asset does not matter because it exists. It matters because it is exposed to a viable attack.

From severity to exploitability. A CVSS of 9.8 only matters if the vulnerability is under active exploitation in the observed context.

From report to mobilization. Each prioritized exposure needs an owner, a deadline and closure validation, with effective mobilization.

THE FIVE PHASES

CTEM is a continuous five-phase cycle

The Gartner framework organizes exposure management into iterative phases, each with a defined output that feeds the next.

Whitepaper 2026
CSURFACE

PHASE 01

Scoping

Definition of the protection scope: surfaces, critical assets, business priorities and regulatory boundaries.

PHASE 02

Discovery

Continuous discovery of assets, vulnerabilities, misconfigurations, identity weaknesses and digital supply chain within the scope.

PHASE 03

Prioritization

Ranking of exposures by probability of exploitation and business impact, rather than by abstract severity.

PHASE 04

Validation

Confirmation through adversarial testing that the exposure is genuinely exploitable, and that remediation closes the attack path.

PHASE 05

Mobilization

Coordinated execution of remediation across security, IT and cloud teams, with an owner, a deadline and closure validation.

The cycle is iterative and recurring

The five phases do not constitute a project with a beginning and an end. They constitute a **permanent cycle**. At each iteration, the Mobilization phase feeds back into Scoping: new business priorities, newly acquired systems, new regulations and new evidence of observed exploitation redefine what will be discovered, prioritized and validated in the next turn.

Gartner describes this iteration as an essential characteristic of the program. Maturity is measured by **cycle speed** and by **coverage across turns**. The faster the cycle operates and the broader the coverage, the greater the likelihood that the program anticipates the threat rather than reacting to it.

What each phase produces for the next.

- **Scoping to Discovery:** a list of what is "in scope" and why.
- **Discovery to Prioritization:** an enriched inventory of exposures.
- **Prioritization to Validation:** a queue of what deserves real testing.
- **Validation to Mobilization:** evidence of exploitability and a remediation plan.
- **Mobilization to Scoping:** lessons learned and adjusted priorities.

Scoping: defining what is protected

The foundational phase of the program. It establishes which assets, processes and surfaces matter to the business, guiding the entire cycle that follows.

What the phase delivers

The Scoping phase answers three questions that guide the entire program.

- **Which assets materialize business risk?** Payment systems, personal data under the CCPA and state privacy laws, intellectual property, critical infrastructure of the sector.
- **Which surfaces will be monitored?** External (public DNS, digital presence, subsidiaries), internal (network, identity, configuration), supply chain (third parties, SaaS, embedded code).
- **Which regulations are in play?** The CCPA and state privacy laws, GLBA and FFIEC guidance, PCI DSS, CIS Controls, sector standards.

Without well-executed Scoping, the program enters Discovery without focus and generates noise instead of risk reduction. The scope must be reviewed each cycle. Newly acquired systems, regulatory changes and M&A activity reorder priorities.

Gartner recommends involving, beyond the security team, representatives from IT, cloud, identity, legal and the business lines, because the scope of the program is a business decision, driven by business leadership.

A typical Scoping scenario. A mid-sized financial institution defines as priority-1 the systems in scope for the GLBA Safeguards Rule and FFIEC guidance, its open-banking APIs, the domains of subsidiaries acquired over the past 5 years and its critical SaaS suppliers. Other assets (intranet, isolated staging environments) remain priority-2, with a slower review cadence.

The common mistake. Treating Scoping as a one-off planning exercise. CTEM requires the scope to evolve each cycle. Assets discovered in the following phase feed back into the definition of what is "in scope" and, therefore, what deserves attention in the next turn.

Discovery: finding what was outside monitoring

The most underestimated phase of CTEM. Discover what actually exists, classify it precisely and assign ownership, before any prioritization.

What the phase delivers

The Discovery phase goes far beyond enumerating known assets. For a mature CTEM program, Discovery covers four dimensions.

- **Assets.** Domains, subdomains, IPs, web applications, APIs, certificates, exposed services, including those of brands and subsidiaries.
- **Vulnerabilities and misconfigurations.** Technical flaws in exposed services, missing headers, expiring certificates, inadequate configurations.
- **Identity.** Leaked credentials, orphaned accounts with external access, granted OAuth scopes, exposed federations.
- **Digital supply chain.** Embedded third-party scripts, CDNs, supplier APIs that execute in the client context.

The output of Discovery is an **inventory enriched with context**: resolved ownership, inferred criticality and cross-referenced exposure signals. A deficient Discovery compromises every subsequent phase of the program.

A typical Discovery scenario. An organization declares 350 external assets in its CMDB. After continuous Discovery, the count reveals 1,200 assets. Only 280 from the original CMDB are valid (some decommissioned, others unmapped). The remaining 920 were discovered through correlation, including subdomains inherited from acquisitions, forgotten staging environments, related brands and assets on cloud providers provisioned by engineering without central registration.

The common mistake. Confusing Discovery with a declarative inventory. Discovery in CTEM does not ask "what do you know you have?". It investigates "what is externally observable and linked to your organization?", including what the organization does not know it has.

Prioritization: prioritizes by real exploitability in the asset's context

Prioritization is where CTEM differs most clearly from traditional vulnerability management. The question moves from "which is the most severe CVE?" to "what is under exploitation now, in my context?".

What the phase delivers

Prioritization integrates three distinct signals to produce the remediation queue.

- **Probability of exploitation.** exploitation-probability index, active-exploitation catalog, evidence of a published exploit, indicators of an active campaign.
- **Business impact.** Criticality of the asset in the business context, sensitivity of the data processed, regulatory exposure.
- **Reachability.** Is the affected asset exposed to the attacker? Is the path open? Are there compensating controls?

The result is a queue that is **radically smaller** than the raw Discovery list. Market analyses indicate that around 2% of exposures lead to critical assets, and more than 75% end in a dead-end path. Correct prioritization is what separates an operable CTEM program from an impossible *backlog*.

Prioritization must be **dynamic**. Changes in the threat landscape, such as a CVE entering the active-exploitation catalog, an exploit being published or a patch being released, must reorder the queue in hours, keeping pace with the threat rather than with monthly cycles.

A typical Prioritization scenario. CVE-2024-XXXX is disclosed with a CVSS of 7.5.

Initially it sits in the "medium" priority band. Three days later, an exploit is published and the CVE enters the active-exploitation catalog. The exploitation-probability index jumps from very low to very high. A vulnerable instance is detected on an asset classified as critical during Scoping. The queue reorders. That 7.5 CVE rises to the top, ahead of several 9.x CVEs with no evidence of active exploitation.

The common mistake. Prioritizing by static CVSS. A CVSS of 9.8 indicates technical severity that is potentially severe in the lab; exploitation in the real context depends on factors of the environment itself. The operational CTEM queue prioritizes by evidence of real exploitability on the asset.

Validation: proving the exposure is real

The phase that separates hypothesis from evidence. Validation answers, through real testing, whether the exposure is exploitable. After remediation, it confirms the path has been closed.

What the phase delivers

Validation is what distinguishes CTEM from an extension of vulnerability management. Instead of assuming that every listed CVE is a risk, the program **empirically tests** whether the exposure is exploitable in the observed context.

Three types of validation feed the CTEM cycle.

- **Technical exploitability validation.** Non-destructive active tests confirm that the vulnerability can be exploited through the declared vector. Includes Breach and Attack Simulation (BAS) and automated pentesting.
- **Reachability validation.** Can the external attacker actually reach the asset? Are there compensating controls (WAF, segmentation, authentication) that neutralize the path?
- **Closure validation.** After remediation, automated retesting confirms that the attack path has been effectively eliminated.

Validation is a rule of quality, applied consistently. Without continuous validation, the program accumulates "remediation" records that do not prove closure.

A typical Validation scenario. An exposure prioritized in the previous phase arrives for validation. The non-destructive active test confirms that the service is vulnerable *and* externally reachable. The team applies a patch and marks it as remediated. Without closure validation, the cycle ends here. With validation, the retest detects that the patch was applied in the staging environment but not in production, and the exposure automatically reopens the ticket.

The common mistake. Limiting validation to annual pentests. The environment changes daily, and an annual pentest delivers outdated evidence. CTEM requires *continuous* and *automated* validation, because what was exploitable yesterday may have been closed by a deployment, and what was secure may have been opened by a configuration change.

Mobilization: from validated alert to delivered fix

The phase where many programs lose traction. Mobilization is the coordination across Security, IT, Cloud and Development so that each validated exposure becomes a delivered fix, with an owner, a deadline and evidence of closure.

What the phase delivers

According to Gartner, "organizations that implement CTEM with strong cross-business mobilization will experience at least a 50% reduction in successful attacks by 2028"^{G2}. The critical point of "cross-business" is that security does not remediate on its own. IT applies the patch, cloud adjusts the configuration, development fixes the code, the supplier updates the component.

Efficient Mobilization depends on three conditions.

- **Identified owner.** Each validated exposure reaches the right team, already linked to the team responsible for the asset.
- **Integration into the existing flow.** The alert enters the Jira, ServiceNow, Slack, Teams or SOAR the team already uses, rather than yet another extra portal.
- **Closure validation.** When the team marks it as "remediated", the cycle confirms that the attack path has been effectively eliminated (Validation feeds back into Mobilization).

Without these conditions, even a technically well-designed CTEM program fails organizationally, because the remediation end is left orphaned.

A typical Mobilization scenario. The CTEM platform produces a validated alert: "CVE-2024-XXXX confirmed exploitable on api-payments.example.com". The alert automatically enters the payments squad's Jira as a critical-severity ticket, linked to the tech lead identified by ownership pattern. The #incident Slack channel is notified. After the patch is deployed, the automated retest confirms closure, and the ticket is closed with attached evidence.

The common mistake. Treating Mobilization as "opening a ticket". Mobilization is orchestration: SLAs by criticality, escalation on inactivity, automatic closure confirmation and MTTR metrics per team. Without this, tickets pile up and the program loses traction.

PUBLIC CASES

How the CTEM paradigm would have contributed

A critical analysis of six public incidents. Each case identifies specifically where the CTEM cycle would have contributed, and where it would not, rather than generalizing.

Whitepaper 2026
CSURFACE

CASE 01 · DECEMBER 2021

Log4Shell (CVE-2021-44228)

A vulnerability in a Java library embedded in thousands of applications. Public exploitation within hours. Organizations with a declarative inventory needed weeks to identify every instance.

Where CTEM would have contributed: Continuous Discovery that maps components embedded at runtime accelerates identification of the reach. Dynamic Prioritization raises it to the top at the moment of disclosure. Validation confirms exploitability per asset, beyond mere binary presence.

Honest limit: no platform would have prevented the event. The gain lies in *reducing the time between disclosure and effective mitigation per asset*.

SOURCE · active-exploitation catalog · Apache

CASE 02 · MAY 2023

MOVEit Transfer (CVE-2023-34362)

SQL injection in a SaaS file-transfer product. More than 2,500 organizations affected through a supplier, frequently without the asset being in the direct inventory. The C10p group exploited it at scale.

Where CTEM would have contributed: Scoping that includes the digital supply chain and contracted SaaS. Discovery maps where the product is in use. Mobilization engages the supplier with a documented SLA, rather than awaiting reactive communication.

Honest limit: the vector was a zero-day in the product. The advantage lies in *reducing the time to organizational awareness* from weeks to hours.

SOURCE · CISA · Progress

CASE 03 · OCTOBER 2023

Citrix Bleed (CVE-2023-4966)

A vulnerability in NetScaler ADC and Gateway allowed session tokens to leak, enabling hijack without credentials. Patches were applied, but pre-existing sessions remained valid. Groups such as LockBit exploited organizations that patched without invalidating sessions.

Where CTEM would have contributed: Continuous Validation verifies the *execution of the session-invalidation procedure*, beyond the presence of the patch. Without that second check, "remediated" indicated only that the patch was installed while the exposure remained open.

Honest limit: capturing this scenario requires Validation that covers *compensating controls*, with verification beyond the mere presence of the patch.

SOURCE · CISA · Mandiant

CASE 04 · JANUARY 2024

Ivanti Connect Secure (CVE-2024-21887 e CVE-2023-46805)

A chain of two vulnerabilities in a remote-access product, exploited as zero-day. Patches were only made available roughly two weeks after disclosure. Initial mitigation via XML configuration. Persistence deployed before the patch survived applications without a full follow-up (factory reset).

Where CTEM would have contributed: Discovery identifies exposed instances even before disclosure. After public IoCs, Validation looks for signs of prior compromise. Mobilization coordinates compensating mitigation (segmentation, reinforced MFA) while the patch is developed.

Honest limit: CTEM does not replace a missing patch. The gain is *operational*: faster compensating mitigation and detection of post-patch persistence.

SOURCE · CISA · Voletixity

CASE 05 · FEBRUARY 2024

ScreenConnect (CVE-2024-1709)

An authentication bypass in ConnectWise ScreenConnect allowed the creation of an administrative account without credentials. Mass exploitation by multiple groups within a few days of disclosure. MSPs and organizations using the product on internal endpoints were especially affected.

Where CTEM would have contributed: Discovery identifies exposed ScreenConnect instances. Prioritization raises it to the top at disclosure. Validation confirms the patch was applied and looks for signs of compromise between disclosure and application. For MSP end clients, the digital supply chain is mapped.

Honest limit: the gain here is clear because a patch exists and there is an observable exploitation window.

SOURCE · CISA · ConnectWise

CASE 06 · JUNE 2024

Polyfill.io supply chain

The domain of a widely embedded CDN was compromised after a change of ownership. More than 100,000 legitimate sites began serving malicious code from within their own HTML. The vector was in the digital supply chain at runtime, outside the attack surface that conventional programs monitor.

Where CTEM would have contributed: Discovery in CTEM covers runtime dependencies. Prioritization raises alerts about a change of ownership of a third-party domain. Mobilization removes or replaces the include via deployment.

Honest limit: a CTEM program without the digital supply chain in Discovery would have missed the event entirely.

SOURCE · Sansec · Cloudflare

The six cases illustrate specific, bounded and realistic contributions of the framework. The CTEM cycle reduces the time between disclosure and effective mitigation through continuous discovery, per-asset validation and coordinated mobilization. It remains dependent on the patch for the fix, on other layers for zero-day control, and on a Discovery that covers the digital supply chain and identity to operate well.

Six mistakes that stall a CTEM program

Patterns observed in implementations that stalled. Each pitfall comes with a recommendation on how to avoid it.

PITFALL 01

Treating CTEM as the purchase of a single tool

CTEM is a program, sustained by process and orchestration. No vendor delivers "a complete CTEM platform". **Avoid:** choosing a single tool and declaring the program done. Maturity comes from orchestrating multiple capabilities, some already present in the organization.

PITFALL 02

Program scope equal to the CMDB inventory

The CMDB lists what the organization knows it has. CTEM must confront what the organization does *not* know it has. **Avoid:** Scoping based only on a declarative inventory. Include extended discovery in the first iteration.

PITFALL 03

Prioritization by static CVSS

CVSS does not distinguish "severe in the lab" from "exploitable now in your context". **Avoid:** using CVSS in isolation. Combine it with exploitation-probability indices, evidence of active exploitation and business impact, and update the queue in hours, keeping pace with the threat rather than with the monthly calendar.

PITFALL 04

Validation only via an annual pentest

The environment changes daily. An annual pentest delivers outdated evidence. **Avoid:** treating validation as an event. Implement continuous, automated validation for high-criticality assets.

PITFALL 05

Mobilization without clear ownership

Alerts with no defined owner pile up until they turn into noise. **Avoid:** launching the program without first ensuring an ownership matrix: who is accountable for each asset, who applies the patch, who decides on an exception.

PITFALL 06

Metrics that measure activity rather than risk

"CVEs fixed per month" rewards volume rather than risk reduction. **Avoid:** running the program from an operations dashboard. Adopt MTTR of *validated* exposures, material risk reduction and prioritization accuracy as the primary metrics.

Where the category stands in the adoption cycle

The status of the framework in the market, formal Gartner projections and the characteristics of the organizations that implement it.

HYPE CYCLE 2024

Slope of Enlightenment

CTEM advanced from the "Peak of Inflated Expectations" to a phase of pragmatic adoption^{G4}

ADOPTION BY 2026

~30%

of cybersecurity programs will have CTEM as their adopted operating model^{G1}

CROSS-BUSINESS MOBILIZATION

< 25%

of organizations in 2025 have mature cross-business mobilization^{G2}

Status in the Hype Cycle

Gartner introduced CTEM in the Hype Cycle for Security Operations in 2022. In 2023, it sat at the "Peak of Inflated Expectations". In 2024 and 2025, it transitioned to the "Slope of Enlightenment", the phase in which the category begins to have practical evaluation criteria and real implementation cases.

This positioning means the framework is sufficiently consolidated to be adopted, and the vendor market is reorganizing around it. Vendors of ASM, BAS, control validation and vulnerability management are shifting their messaging toward "supporting CTEM".

Profile of the organizations that adopt it

Market research in 2024 and 2025 indicates that CTEM is adopted predominantly in organizations that have at least one of the characteristics below.

- **Strong sector regulation** (finance, healthcare, critical infrastructure).
- **Prior maturity in traditional VM** that has reached its effectiveness ceiling.
- **Large, dynamic digital surfaces** (e-commerce, digital banking, SaaS).
- A history of a **material incident** that exposed the fragility of the point-in-time model.

Consolidated Gartner projections.

- **2026:** organizations with CTEM are 3× *less likely* to suffer breaches.
- **2028:** CTEM with cross-business mobilization reduces successful attacks by ≥50%.
- **2028:** >50% of critical exposures will stem from non-technical origins: identity, SaaS, third parties, process.

Progressive maturation is the norm. Effective CTEM programs typically take between 12 and 24 months to reach the Coordinated stage. Implementations at early stages (Reactive, Reconnaissance) still deliver measurable gains within a few cycles. The program does not require an initial big-bang. It advances through consolidation at each iteration.

^{G4} Gartner, "Hype Cycle for Security Operations, 2024".

Four stages of CTEM maturation

The maturity of a CTEM program is not measured by a checklist. It is measured by cycle speed, by coverage across turns and by the proportion of the program that operates in an automated way.

STAGE 01

Reactive

Monthly point-in-time scan, CVE list ranked by CVSS, remediation on demand. Partial coverage of the external surface, without visibility of identity or supply chain.

STAGE 02

Reconnaissance

External inventory discovered continuously (operational EASM). Prioritization incorporates exploitation-probability indices or active-exploitation catalogs. Point-in-time validation via annual pentest. Mobilization via manual tickets.

STAGE 03

Coordinated

Coverage extended to identity and the digital supply chain. Dynamic prioritization with proprietary threat intel. Automated exploitability validation. Mobilization integrated with SIEM, ticketing and SOAR.

STAGE 04

Operational

Complete CTEM cycle in <24h. Continuous closure validation. Financial quantification of risk. Formalized cross-business mobilization. Measurable risk reduction each cycle.

What differentiates each stage

Maturation is evolution through consolidation. **Stages are not skipped:** each requires the previous one to be stable. Running continuous Validation over a deficient Discovery generates noise. Implementing dynamic Prioritization without proprietary threat intel delivers indefensible rankings. Organizations progress from Stage 02 to 03 by investing in *additional coverage* (identity, supply chain) and *proprietary threat intel*. They progress from 03 to 04 by investing in *cycle speed*, *continuous validation* and *financial quantification*.

Common problems in tools used in the program

PARTIAL COVERAGE

EASM limited to subdomain enumeration

Platforms that start only from the domains already declared by the client, enumerate subdomains and inspect the SAN in certificates. **Limit:** they do not discover subsidiaries, related brands, cloud shadow IT or assets with no direct DNS link to the root domain.

INCOMPLETE VISIBILITY

Digital supply chain and identity out of scope

Tools focused on technical vulnerability that do not cover embedded scripts and CDNs, third-party APIs, leaked credentials or federation exposures. **Limit:** more than 50% of the critical exposures forecast for 2028 will stem from non-technical origins.

POOR INTEGRATION

Platforms that do not export data

Tools with their own dashboard, no documented API, no webhook, no native connector for SIEM, ticketing or SOAR. **Limit:** it makes automated Mobilization impossible and forces the team to operate in an additional portal, instead of the existing flow.

MISSING CONNECTORS

No integration with the installed stack

Platforms that do not cover the environment-specific connectors: enterprise SIEM (Splunk, Sentinel, QRadar, Elastic), ticketing (Jira, ServiceNow), SOAR (Tines, XSOAR), GRC. **Limit:** the team gains alerts but cannot orchestrate response without manual work.

SILOED VIEW

Tools that do not talk to one another

EASM, BAS, vulnerability management, credential monitoring and GRC operate in silos, with independent dashboards. **Limit:** without consolidation, the CTEM cycle has to be reassembled manually by the team, phase by phase, losing speed and introducing discrepancies.

UNDEFINED OWNERSHIP

Discovery without reliable attribution

Platforms that discover assets but do not resolve whom they belong to within the organization. **Limit:** the team spends capacity on manual ownership triage before it can prioritize or mobilize. A false positive in ownership undermines trust across the whole cycle.

What is measured in a CTEM program that works

CTEM metrics are intentionally different from vulnerability-management metrics. The focus is on demonstrating a reduction of material exposure over time, beyond the mere count of fixed CVEs.

METRIC	WHAT IT MEASURES	CADENCE
Critical-surface coverage	Percentage of the critical assets defined in Scoping that are under continuous discovery.	Weekly
MTTR of validated exposures	Average time between exploitability validation and confirmation of closure. It replaces the MTTR of "open CVEs".	Monthly
Prioritization accuracy	Proportion of alerts that support concrete action (measured by the rate of real closure vs. dismissal as false positive).	Monthly
CTEM cycle speed	Time between the publication of a relevant new threat and the program's adaptation (revalidation, reprioritization, mobilization).	Per event
Open critical exposures	Absolute volume and trend of <i>validated</i> and <i>unremediated</i> exposures on critical assets.	Weekly
Material risk reduction	Financial quantification (FAIR, VaR) of aggregate exposure over time. The program's ultimate metric.	Quarterly
Rediscovery rate	Proportion of discovered assets that reappear in subsequent cycles after remediation. It indicates configuration drift or reappearance.	Monthly

What should *not* be a primary metric. The number of open CVEs, the average severity of the backlog, or the volume of patches applied. These are metrics of internal operation, distinct from risk reduction, and they tend to reward activity rather than outcome. A mature CTEM program has *few* metrics, all tied to observable reduction of material exposure.

How CSURFACE contributes to your CTEM program

A direct phase-by-phase mapping, and how CSURFACE integrates with the existing SIEM, SOAR, ticketing and GRC stack already in operation in the organization.

Whitepaper 2026
CSURFACE

CTEM PHASE	WHAT THE PHASE REQUIRES	HOW CSURFACE CONTRIBUTES
PHASE 01 Scoping	Definition of critical assets, surface scope and regulatory priorities.	Discovery with no input reveals the real scope (subsidiaries, brands). Contextual classification categorizes each asset. Tracks: CCPA, GLBA/FFIEC, CIS, PCI DSS.
PHASE 02 Discovery	Enriched inventory covering assets, vulnerabilities, identity and the digital supply chain, with resolved ownership.	Machine Learning and an agentic layer with high precision. Agentless. Digital supply chain (scripts, CDN, API) and validated leaked credentials.
PHASE 03 Prioritization	Dynamic ranking by exploitability, impact and reachability.	A proprietary Threat Sensor correlates NVD, exploitation-probability indices, active-exploitation catalogs, active exploits. Queue recalculated in <24h after a signal change.
PHASE 04 Validation	Empirical confirmation of exploitability and of closure, continuously where test coverage exists, with passive monitoring across the rest of the surface.	Non-destructive validation , active exploitability validation where test coverage exists, with passive monitoring across the rest of the surface. Continuous Validation after patching. Active validation of credentials before the alert.
PHASE 05 Mobilization	Distribution to the owner, integration with the existing flow, closure validation.	Native integrations : SIEM, ticketing, ChatOps, SOAR. REST API and Webhooks for custom integration.
CROSS. Financial quantification	Translation of technical risk into financial value for the board.	CRQ with FAIR , IBM 2025 United States, VaR P90 and P99, regulatory exposure. Calibration per asset and process for clients.

How CSURFACE connects to the stack you already have

EXISTING SIEM Splunk · Sentinel · QRadar · Elastic Validated CSURFACE alerts enter the SIEM already prioritized and with ownership resolved, eliminating the noise of manual correlation over unvalidated findings.	EXISTING SOAR Splunk SOAR · Tines · Cortex XSOAR CSURFACE webhooks trigger playbooks: notify the squad, open a ticket, isolate a host, escalate on inactivity. Orchestration without reimplementing logic.
EXISTING TICKETING Jira · ServiceNow Each validated exposure becomes a ticket in the existing flow, linked to the squad responsible for the asset. Automatic closure once Continuous Validation confirms remediation.	EXISTING VM SCANNER · GRC Tenable · Qualys · Rapid7 · Archer CSURFACE complements the internal scanner with external coverage, real exploitability validation and digital supply chain context. For GRC, it feeds auditable evidence of exposure reduction.

CSURFACE complements the existing stack rather than replacing it. It acts as a layer of discovery, prioritization and validation that connects to the platforms already in operation, adding capabilities those platforms do not cover natively.

Consolidated bibliography: primary and complementary sources

A list of the references used in this whitepaper, organized by category. Recommended material for further reading.

Gartner, primary sources

- **Gartner (2022).** "Implement a Continuous Threat Exposure Management (CTEM) Program". The document that formalized the CTEM concept and described the five phases.
- **Gartner (2023).** "Top Trends in Cybersecurity for 2023". Positioning of CTEM among the strategic trends of the year.
- **Gartner (2024).** "Hype Cycle for Security Operations". Positioning of CTEM on the Slope of Enlightenment.
- **Gartner (2024).** "How to Manage Cybersecurity Threats, Not Episodes". The argument for the continuous-program framework.
- **Gartner, 2026 projection.** Organizations with CTEM are 3× less likely to suffer a breach.
- **Gartner, 2028 projection.** CTEM with cross-business mobilization reduces successful attacks by ≥50%.
- **Gartner, 2028 projection.** >50% of critical exposures will stem from non-technical origins.

Complementary technical sources

- **Active-exploitation catalog.** A list of vulnerabilities demonstrably exploited, maintained by public exploitation-intelligence sources.
- **Exploitation-probability index.** A probabilistic model of exploitation over the next 30 days, drawn from public exploitation-intelligence sources.
- **Open FAIR (Open Group).** A framework for the probabilistic quantification of cyber risk. opengroup.org/openfair
- **IBM Cost of a Data Breach Report 2025.** A study conducted by the Ponemon Institute. ibm.com/reports/data-breach
- **Sansec (2024).** "Polyfill Supply Chain Attack". A technical analysis of the June 2024 incident.
- **Mandiant (2023).** "Citrix Bleed CVE-2023-4966". Technical analysis and indicators.
- **Voletixity (2024).** "Ivanti Connect Secure VPN Compromise". The initial discovery of active exploitation.
- **CISA.** Official advisories on Log4Shell, MOVEit Transfer, Citrix Bleed, Ivanti Connect Secure and ScreenConnect.

This whitepaper consolidates public presentations of the CTEM framework by Gartner, complemented by technical documentation from CISA and the Open Group. It is not a document sponsored by any vendor other than CSURFACE as publisher, and it seeks to represent the framework neutrally before describing the specific contribution of CSURFACE.

NEXT STEPS

CTEM is a continuous program. Start with the **diagnosis of your surface.**

A mature CTEM program depends on reliable discovery, defensible prioritization and empirical validation. CSURFACE offers three entry points, all without commercial friction, so the organization can assess the technical contribution before any commitment.

1. Free preliminary analysis

Provide the company domain and receive, within 48h, the map of external exposure from an attacker's point of view. Subdomains, S3 buckets, leaked credentials, third-party scripts and forgotten portals. No card. No meeting.

2. Public risk calculator

Estimate, in US dollars, the expected annual loss from an incident, using the FAIR methodology, the IBM 2025 United States benchmark and VaR P90 and P99. Material for the audit committee and the board. No sign-up.

3. Technical architecture conversation

When the organization is evaluating tools to support CTEM, a one-hour conversation on architecture is worthwhile: Machine Learning, the agentic layer, integrations and roadmap. No commercial pressure.

See your company's **external exposure in 48h**

[Get the free analysis](#)