



HORIZON3

AI Changed the Economics of Cyberattacks: Why Security Needs a New Exposure Management Operating Model

The future of cybersecurity won't be defined by who sees the most. It will be defined by who can prove the most.

Executive Summary

Every generation of cybersecurity has been built on a set of assumptions. AI just invalidated them.

For decades, security programs were built around a simple premise: attackers operate at human speed. That assumption shaped everything from annual penetration tests and vulnerability management to incident response, compliance programs, and security operations.

Artificial intelligence changed the premise.

Today's attackers can discover exposures, develop exploits, and identify attack paths at machine speed. Activities that once required days or weeks for skilled teams to accomplish can increasingly be compressed into minutes and repeated at unprecedented scale.

The result isn't a new attack surface. It's a fundamentally different operating environment for exploiting the one organizations already have.

Organizations cannot respond by doing more of the same: more vulnerability scans, more dashboards, and more alerts. Those approaches improve visibility, but that alone does not answer the question that now matters most:

→ **Can attackers actually compromise my environment today?**

That question requires a new security operating model built around continuous exposure management, informed action, and verified outcomes. CTEM provides a framework for operationalizing that model by continuously discovering, validating, prioritizing, and reducing exposures as environments change. At its core, CTEM brings continuous discipline to exposure management, helping organizations focus on the exposures most likely to create meaningful business risk.

Every generation of cybersecurity has been built on a set of assumptions. AI just invalidated them.

AI Didn't Create the New Attack Surface. It Changed the Rules.



For decades, cybersecurity optimized for visibility using this traditional security model:

Discover → Identify → Protect → Monitor

It was built on a simple set of assumptions: attackers operated at human speed, periodic assessments provided an accurate picture of risk, visibility created confidence, and compliance demonstrated readiness.

However, security leaders are now asking fundamentally different questions.

Yesterday's Question	Today's Question
How many vulnerabilities do we have?	Which ones can actually be exploited today?
Are our security controls deployed?	Can we prove they stop real attacks?
Are we compliant?	Are we resilient against machine-speed attackers?

These questions require more than visibility. They require evidence.

The assumptions that shaped modern cybersecurity no longer hold. Organizations need a fundamentally different security operating model, one built for the realities of machine-speed attackers rather than the assumptions of the past.

The rest of this paper explains why those assumptions no longer hold, how AI changed the economics of cyberattack, and what organizations must do differently.

Those questions require more than visibility.
They require evidence.

The Attack Surface Grew. AI Accelerated **How Attackers** **Exploit It.**



The attack surface has been expanding for decades. AI didn't start that trend.

Cloud computing expanded environments beyond traditional data centers. SaaS applications multiplied. APIs connected systems that were once isolated. Identity became the new perimeter. Third-party software became integral to everyday operations. Organizations are now deploying AI infrastructure alongside everything else they already manage.

What is new is the speed at which attackers can understand and exploit that complexity.

Complexity was manageable when attackers operated at human speed. Artificial intelligence changed that equation by dramatically reducing the time and effort required to analyze complex environments and uncover opportunities that would have been impractical to find manually.

What Changed: Complexity and the Speed of Exploitation

Yesterday	Today
More assets	More relationships between assets
Servers and endpoints	Cloud, SaaS, APIs, identities, third parties, AI infrastructure
Individual vulnerabilities	Interconnected systems
Attackers analyzed complexity manually	AI analyzes complexity at machine speed

That's the shift.

The challenge is no longer the number of assets organizations must defend. It's the growing gap between the speed at which attackers can exploit exposures and the speed at which security teams can understand and reduce them.

Modern attack surfaces contain not only more assets, but far more relationships among assets, identities, applications, cloud services, APIs, and third-party systems.

AI Changed the Economics of Cyberattack



For decades, defenders benefited from one simple constraint: Every stage of a cyberattack required skill, time, and effort.

Reconnaissance took days. Vulnerability research took weeks. Developing reliable exploits required specialized expertise. Every campaign consumed finite resources, forcing attackers to make deliberate choices about where to invest their time. **Artificial intelligence changed those economics.**

AI dramatically reduces the cost of understanding complex environments, identifying exploitable weaknesses, discovering attack paths, and adapting as conditions change. Work that once required teams of experienced operators can increasingly be performed by a single individual working with AI.

How AI Changed the Economics of Cyberattack

Before AI	With AI
Reconnaissance took days	Reconnaissance happens in minutes
Skilled researchers analyzed targets	AI analyzes environments at machine speed
Attack paths were built manually	AI identifies attack paths automatically
Campaigns required teams	One operator performs the work of many
Attackers optimized for effort	Attackers optimize for scale

Security programs were built around attackers having limited time and resources. Artificial intelligence removed that constraint. When attacks become dramatically cheaper, more attacks become economically worthwhile.

Organizations that were once too difficult or too expensive to target become viable. More attack paths become worth pursuing. More campaigns become economically worthwhile. The result isn't simply greater attacker productivity. It is a fundamental shift in the economics of cyberattack.

Visibility Solved Yesterday's Problem.



For years, cybersecurity's hardest problem was understanding the environment:

- What assets were connected?
- What software was running?
- What vulnerabilities were present?

Organizations invested heavily in asset inventories, vulnerability scanners, threat intelligence, and security telemetry because visibility led to better decisions.

That operating model worked because understanding the environment was the hardest part.

Today, understanding the environment is no longer enough.

Security teams have never had more information. They have more dashboards, alerts, vulnerability data, and telemetry than at any point in the industry's history. Yet many organizations remain uncertain about a far more important question.

Can attackers actually compromise my environment today?

Visibility Answers...	Evidence Answers...
What assets exist?	What can attackers actually exploit?
What vulnerabilities are present?	Which exposures create meaningful risk?
What security controls are deployed?	Do those controls stop real attacks?
Was remediation completed?	Did remediation eliminate exploitable risk?

Visibility can tell you where vulnerabilities exist. It cannot tell you which ones attackers can actually exploit, how they might be combined to reach critical systems, whether security controls will stop a real attack, or whether remediation actually reduced meaningful risk.

That's why cybersecurity is no longer an information problem. It's increasingly an evidence problem.

Organizations don't need more information. They need proof.

- Proof that security controls work as intended.
- Proof that remediation reduced meaningful risk.
- Proof that the tested attack paths no longer produce their demonstrated impact.

Visibility tells you what exists. Evidence tells you what matters.

That's the difference between knowing your environment and understanding your risk.

Visibility tells you what exists.
Evidence tells you what matters.

A New Security Operating Model: Exposure Management

Every major shift in cybersecurity has required a new operating model.

Perimeter security gave way to detection and response. Vulnerability management evolved into exposure management. **CTEM made that discipline continuous.**

Artificial intelligence accelerates the need for the next evolution.

The question is no longer whether organizations understand their environment. It's whether they can continuously adapt as quickly as attackers do.

That requires making exposure management a continuous operating discipline, not another periodic security activity.

Questions That Define a Modern Security Operating Model

Yesterday	Today
What do we have?	What can actually be exploited?
What should we fix?	What creates the greatest business risk?
Did we complete the work?	Did our actions reduce meaningful risk?
When is our next assessment?	What changed since our last assessment?

Those questions fundamentally change what security programs optimize for.

Instead of optimizing for visibility, security programs optimize for outcomes. Success is no longer measured by the number of vulnerabilities identified or the volume of alerts investigated. It's measured by whether defensive actions meaningfully reduce exploitable risk and improve the organization's ability to withstand attack.

CTEM provides the framework for operationalizing continuous exposure management by continuously discovering, validating, prioritizing, and reducing exposures over time.


The model works if organizations can execute it continuously.

Continuous Exposure Validation Closes the Loop

A modern security operating model only works if organizations can continuously measure whether their defensive actions achieve the intended outcome.

Continuous validation makes that possible. It provides the evidence needed to determine what attackers can actually exploit, guides defensive decisions, and verifies whether those actions meaningfully reduced risk.

Continuous Exposure Validation Drives Continuous Improvement



Exposure validation becomes the mechanism that drives continuous improvement.

A validated exposure might lead to precision patching. Another may require tuning detections and response, deploying deception, proactively hunting for threats, or reducing blast radius. The appropriate response depends on the exposure, the organization's objectives, and the desired defensive outcome.

Security improves when every decision begins with evidence and ends with verification. That transforms security from a collection of disconnected activities into a continuous operating discipline. Each cycle produces better evidence, better decisions, and stronger defenses, continuously improving the organization's ability to withstand attack.

Cyber Resilience Has a New Standard



For years, cyber resilience was measured primarily by an organization’s ability to recover.

- How quickly could critical systems be restored?
- How effectively could operations continue?
- How rapidly could incident response contain an attack?

Those capabilities remain essential. But recovery begins after an attacker has succeeded. Artificial intelligence raises the standard for resilience.

The New Standard for Cyber Resilience

Yesterday	Today
Recover from attacks	Withstand attacks
Measure response	Measure outcomes
Restore operations	Continuously reduce exploitable risk
Assume controls work	Prove controls work
React	Continuously improve

Resilience can no longer be measured solely by how well organizations recover from attacks. It must also be measured by how effectively they prevent attackers from achieving their objectives and how confidently they can demonstrate their defenses are working.

That changes how organizations measure resilience.


Resilience is not only measured by how well you recover. It’s measured by how confidently you can demonstrate your defenses withstand the attacks that matter most.

In the age of machine-speed attackers, resilience isn't a destination. It's a continuous process of measuring outcomes, adapting defenses, verifying results, and improving over time.

Organizations that embrace that mindset won't simply respond faster. They'll make better decisions, strengthen their defenses, and continuously increase their ability to withstand attack.

That's the new standard for cyber resilience.

Resilience is no longer measured by how well you recover. It's measured by how confidently you can demonstrate your defenses withstand the attacks that matter most.



Conclusion

Every generation of cybersecurity has been built on a set of assumptions.

For decades, those assumptions held. Attackers operated at human speed. Periodic assessments were sufficient. Visibility created confidence. Security programs were designed for an environment where defenders had time to understand, prioritize, and respond.

Artificial intelligence invalidated those assumptions.

The challenge organizations face today isn't a new attack surface or more sophisticated attackers. It's that the economics of cyberattacks have fundamentally changed. Attackers can identify opportunities, adapt their techniques, and scale their operations faster than security programs designed for a different era can respond.

The answer isn't more tools or more information.

It's a new security operating model built around continuous evidence, informed action, and verified outcomes.

Organizations won't succeed by trying to predict every attack. They'll succeed by continuously producing the evidence needed to make better decisions, strengthen their defenses, verify meaningful risk has been reduced, and adapt as quickly as attackers do.

Put Continuous Exposure Management Into Practice

The CTEM Playbook shows how to turn this operating model into a repeatable cycle: discover exposures, validate what attackers can exploit, prioritize by impact, remediate, and verify that fixes worked.

[Explore the CTEM Playbook](#) [Want more information on CTEM?](#)