cymulate

Agentic Cyber Defense Engineering

Beyond Mythos: A New Operating Model for Proactive Security

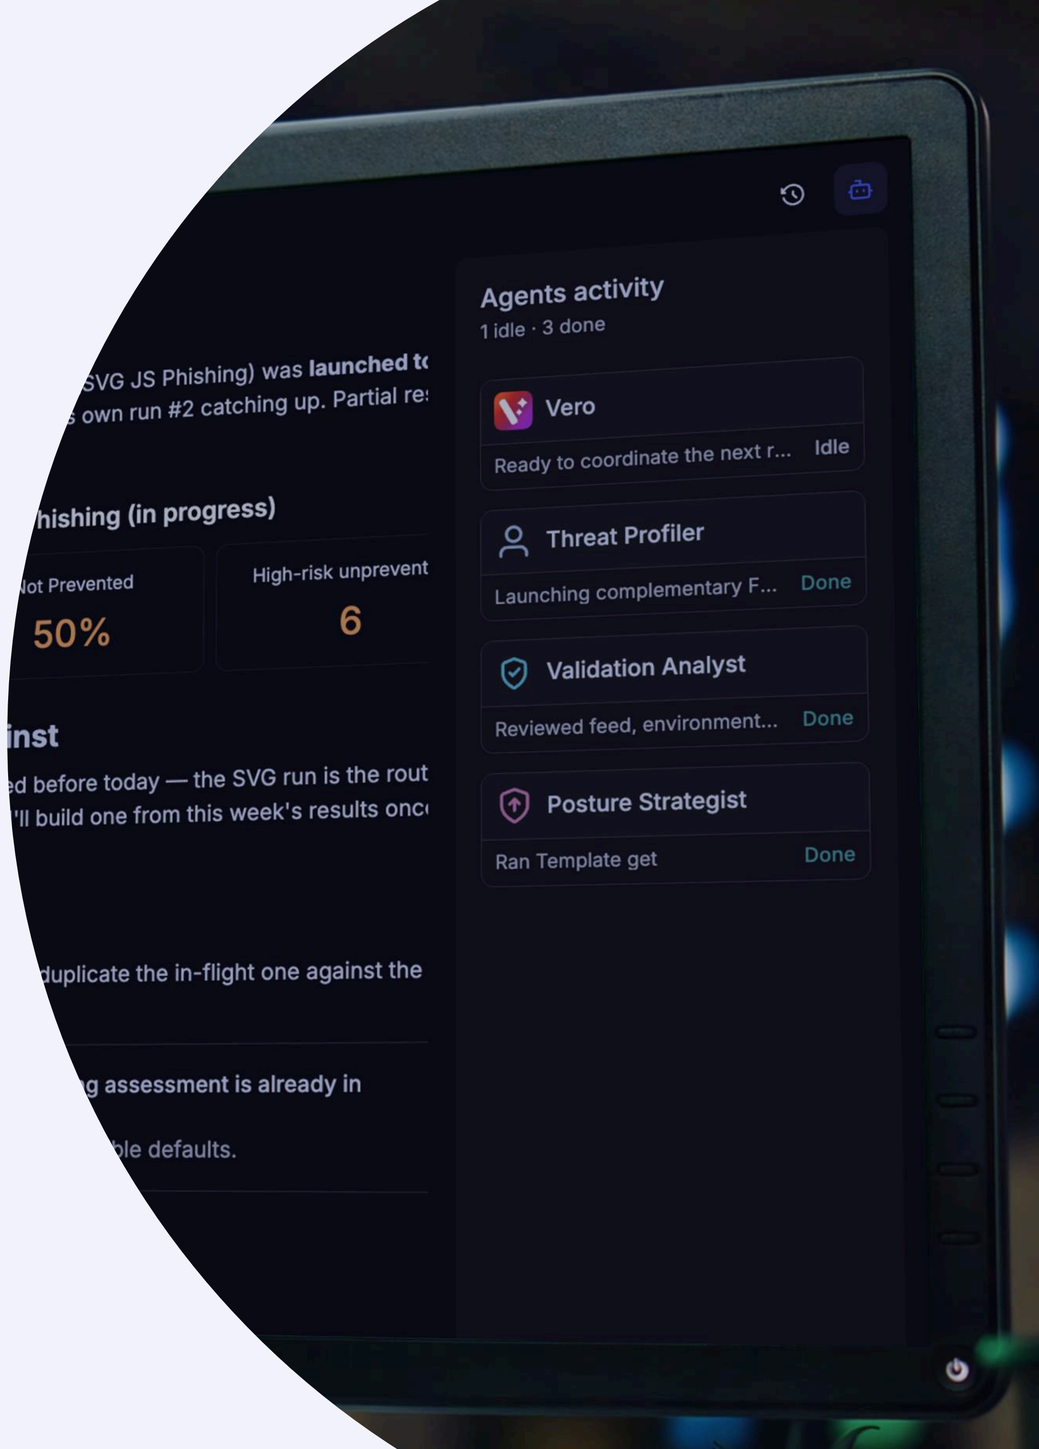

Comprehensive Attack Simulations Across Multiple Threat Vectors

Security teams are tasked with defending against evolving threats while managing 40+ disconnected tools with limited resources. The result is fragmented operations, increased complexity and visibility gaps.

The solution isn't more tools; it's integrated intelligence that unifies security operations, automates workflows and continuously strengthens cyber resilience.

Key Challenges

Tool Sprawl

40+ disconnected security solutions with overlapping capabilities.

Siloed Operations

Disjointed processes across security, IT and risk teams.

Lack of Integration

Limited interoperability with gaps in visibility and response.

Manual Effort

Time-consuming, reactive workflows instead of automated, continuous validation.

Unproven Resilience

Lack of proof that controls deliver resilience against evolving threats.

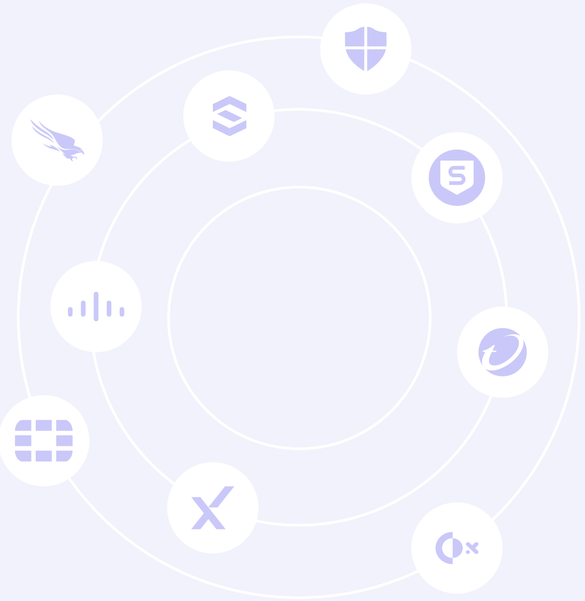

As a result, organizations struggle day-to-day with the reality.

Despite significant investments in security, many organizations still struggle to effectively prioritize and remediate risk. Existing tools and processes often provide visibility into risk and exposures but fall short in driving timely action.

Every exposure represents an opportunity for attackers. To stay ahead of evolving threats, organizations need more than visibility. They need a solution that continuously identifies, validates, prioritizes and remediates risk before exposures can be exploited and threats become incidents.

43%

Of enterprises lose existing customers after a **breach**

Source: Cyber Magazines

32%

of SecOps teams **say they have too many exposures** to prioritize

Source: Gartner

6+

days to recover from a cyber incident

Source: Cyber Magazines

Human-centric security models fail in an AI world

Even as security models have evolved from reactive to proactive, their speed and effectiveness have not kept pace with the reality of post-Mythos threats. Frontier AI models like Mythos give unprecedented speed and scale to exposure discovery and weaponization. This significantly lowers the time-to-exploit from weeks to hours and minutes.

While proactive approaches emphasize broader discovery to cover more of the attack surface, security teams face a tidal wave of both urgency and noise. Human-driven security operations cannot keep up with the scale, speed and sophistication of modern threats, especially those augmented by artificial intelligence. The next wave of frontier AI models will enable attackers to discover, exploit, adapt, and scale attacks autonomously at machine speed.

Agentic Cyber Defense Engineering

Agentic cyber defense engineering is a closed-loop approach that uses AI agents and continuous validation to help security teams prove what works, prioritize what matters and adapt defenses faster than attackers can exploit gaps.

Powered by Vero AI, Cymulate delivers agentic cyber defense engineering thats transform security to autonomous process to continuously prove and improve cyber defenses, so you stay ahead of threats unique to your environment.

The result: A continuously improving security posture that adapts in real time, automatically closes security engineering gaps and drastically reduces exposure windows.

Closed-Loop Process to Continuously Prove and Improve

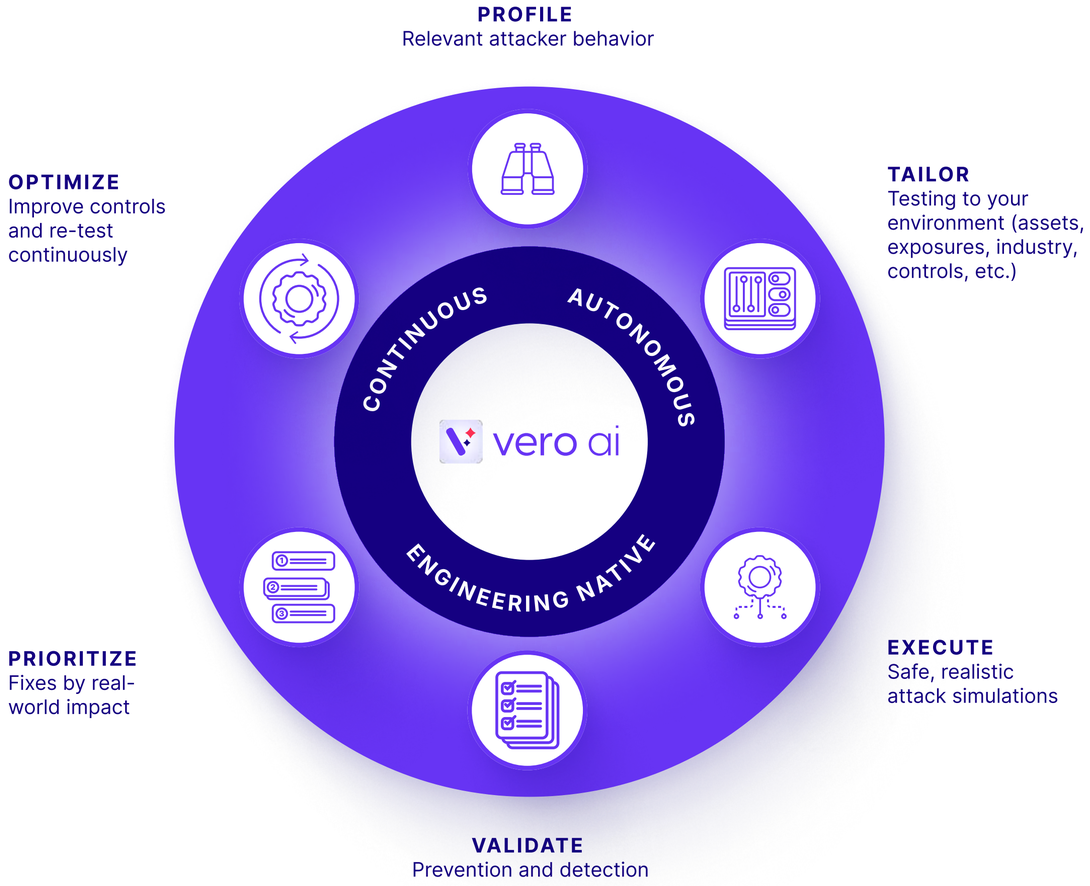

Profile

Tailor

Execute

Validate

Prioritize

Optimize

Agentic Cyber Defense Engineering: Profile

The Problem: Most security programs operate with an outdated understanding of their attack surface. New assets are deployed, controls change, vulnerabilities emerge and threat actors continuously evolve their tactics. Security teams are often forced to make decisions using incomplete or stale information, creating blind spots that attackers exploit.

Traditional security validation relies on periodic assessments and manually curated threat intelligence, leaving organizations exposed to threats that emerge between testing cycles.

The Agentic Solution: The profile phase continuously maintains a real-time understanding of the organization's environment, security controls, attack surface and threat landscape by:

- Modeling and maintaining a profile for the environment and security architecture
- Identifying critical assets, exposures and security controls
- Continuously ingesting new threat intelligence
- Monitoring for emerging attacker techniques and campaigns
- Determining which threats are most relevant to the organization

The Result: Rather than requiring analysts to manually determine what should be tested, the platform continuously understands what matters and updates its threat model accordingly.

Real-World Example

When a new ransomware campaign began targeting manufacturers, Cymulate agentic cyber defense engineering automatically prioritized the threat and added it to the organization's testing profile to validate defenses and close gaps before attackers could strike.

Agentic Cyber Defense Engineering: Tailor

The Problem: Generic security testing often produces irrelevant results. Organizations waste time validating attack techniques that do not apply to their environment while overlooking scenarios that pose actual business risk.

Security teams spend time manually determining which attack paths, assets, integrations, and variables should be included in each assessment. This time-consuming process that is difficult to scale fails in an AI-centric world.

The Agentic Solution: The tailor phase automatically converts environmental context and threat intelligence into customized attack assessments by:

- Matching new threats to relevant assets and controls
- Selecting applicable attack techniques and scenario
- Automatically configuring assessments variables such as environment, test points and integrations

The Result: Every assessment becomes specific to the organization's actual environment rather than relying on generic attack simulations and security teams save significant time on configuring assessments.

Real-World Example

When a new Microsoft 365 credential theft campaign emerged, Cymulate agentic cyber defense engineering automatically generated a tailored assessment to validate the organization's identity, email, endpoint and detection defenses. The financial institution quickly identified security control gaps and strengthened resilience against active threats.

Agentic Cyber Defense Engineering: Execute

The Problem: Security control and threat validation is complex, resource-intensive and requires specialized expertise. Most organizations either outsource validation or rely on costly internal red teams.

As modern threats evolve daily, periodic validation is not enough. Organizations need continuous validation to ensure their security controls remain effective, provide comprehensive coverage, and can detect and prevent the latest attack techniques.

The Agentic Solution: The execute phase automatically launches relevant attack simulations in a safe and controlled manner by:

- Executing configured assessments automatically
- Monitors for triggers on when to re-test based on auto-fixes and changes to environments
- Runs attack actions without impacting operations
- Cleans up activity post testing

The Result: Organizations gain the confidence that their defenses are continuously validated against the threats that matter most. By automatically testing the right controls at the right time, security teams can ensure coverage keeps pace with an evolving threat landscape while minimizing operational overhead and risk.

Real-World Example

The attack begins by stealing credentials from a compromised system, then uses those credentials to move laterally and remotely access additional systems. Once access is established, a malicious payload is downloaded and executed to further compromise the environment. If the attack was successful, the downloaded malicious payload is removed.

Profile

Tailor

Execute

Validate

Prioritize

Optimize

Agentic Cyber Defense Engineering: Validate

The Problem: Stopping an attack is only half the challenge. Security teams must also know whether controls prevented the attack and whether security operations teams would have detected malicious activity if prevention failed.

Many organizations struggle to correlate attack activity with actual control performance, leaving uncertainty about security effectiveness.

The Agentic Solution: The validate phase automatically determines whether attacks were prevented and detected by:

- Correlating attack activity with security controls
- Determining whether attacks were blocked
- Analyzing logs, alerts and telemetry automatically with integrations
- Cleans up activity post testing

The Result: Security teams gain measurable proof of security effectiveness rather than assumptions and identification of all security gaps against the latest threats.

Real-World Example

A malicious PowerShell command is executed on an endpoint to simulate an attacker downloading and launching a payload. The platform validates whether endpoint security blocked the execution, whether PowerShell activity was detected and whether the EDR and SIEM generated the expected alerts.

Profile

Tailor

Execute

Validate

Prioritize

Optimize

Agentic Cyber Defense Engineering: Prioritize

The Problem: Security teams are overwhelmed by a growing volume of vulnerabilities, misconfigurations and attack surface exposures, making it difficult to determine which issues represent real business risk.

Traditional prioritization approaches rely heavily on CVSS scores, threat feeds and static vulnerability data, often lacking the context needed to understand actual exploitability within the environment.

The Agentic Solution: The prioritize phase correlates ingested vulnerabilities and validation results with environmental context to identify the most important exposures by:

- Correlating findings across controls and attack paths
- Identifying exposures that can actually be exploited
- Calculating true risk scores based on business context
- Prioritizing remediation based on real-world impact

The Result: Organizations efficiently allocate resources by focusing on the exposures that attackers can truly leverage

Real-World Example

Two vulnerabilities receive critical CVSS scores. After validation, only one is proven exploitable through an attack path that bypasses existing controls. The platform automatically elevates this finding for immediate remediation while reducing focus on lower-risk issues.

Profile

Tailor

Execute

Validate

Prioritize

Optimize

Agentic Cyber Defense Engineering: Optimize

The Problem: Most organizations can identify security gaps but struggle to close them quickly. Security teams must manually research indicators, write detections, update controls and validate fixes, which often takes weeks or months.

During that time, attackers can continue exploiting known weaknesses.

The Agentic Solution: The optimize phase automatically improves security controls based on validated findings by:

- Generating remediation recommendations automatically
- Auto-pushing IOC updates directly into security controls
- Building new detection rules and auto-deploying rules through integrations
- Re-test controls continuously to confirm effectiveness

The Result: Organizations transform validation into a closed-loop optimization process that continuously adapts and improves threat resilience.

Real-World Example

A detection gap is identified during a ransomware assessment. The platform automatically generates a new detection rule, pushes it to the SIEM and re-runs the assessment to confirm that the threat is now detected. The new rule is automatically ingested into inventory of rules, mapped to the attack library and automated testing is schedule. Any update to the rule triggers a new test.

Agentic Cyber Defense Engineering Control Plane

How it Works:
The Cymulate platform includes an agentic cyber defense engineering plane. Powered with Cymulate Vero AI, the control plane provides an autonomous closed-loop with triggers such as new threat intelligence, security control changes and assessment findings. The independent, autonomous agents automate manual security tasks, such as identifying relevant threats, mapping attack scenarios, executing attacks and automating mitigation.

These agents feed the control plane, which continuously tailors, executes, validates, prioritizes and tunes security assessments. Through integrations with security controls across the entire architecture, the system continuously improves cyber defenses and strengthens threat resilience at machine speed.

Use Cases

Agentic cyber defense engineering unifies security operations, detection engineering, security engineering and compliance into a continuous, automated defense cycle. It continuously validates defenses against the latest attacker techniques, identifies gaps across security controls and drives remediation before threats become incidents.

As new threats and zero-days emerge, the system automatically assesses exposure, tests defenses and adapts protections in real time. The result is continuous assurance, measurable resilience, and always-on proof that security controls are working as intended.

Zero Day & Emerging Threats

Instantly validates exposure to new threats relevant to your environment.

Detection Optimization

Continuously ingest existing detection rules and fixes identified broken rules or gaps.

Fine-Tune Security Controls

Automatically runs attacks and auto-pushes fixes to controls to mitigate threats.

Vulnerability Validation & Prioritization

Continuously ingests assets and vulnerabilities, maps to attacks and calculates actual risk score.

Compliance & Audit Automation

Provides continuous proof of control effectiveness mapped to global common frameworks.

Bottom Line

As cyber threats continue to evolve, security teams can no longer rely on periodic testing, manual processes and disconnected workflows to measure resilience. Agentic cyber defense engineering transforms security validation into a continuous, autonomous capability, ensuring defenses are always aligned to the latest threats, exposures, and business risks.

Ready to move from reactive security operations to continuous cyber resilience?

Learn how Cymulate agentic cyber defense engineering can help your organization prove, prioritize and adapt at machine speed.

See Cymulate Vero AI

Discover Cymulate Vero AI through a quick, interactive guide of our platform **See Vero AI in Action**

Schedule a Demo

Learn how you can start using Agentic Cyber Defense Engineering **Request a Demo**

About Cymulate

Cymulate is the leader in proactive, AI-powered security that continuously proves, prioritizes and adapts against real attacker behavior – before incidents occur. More than 1,000 enterprise security teams rely on Cymulate for autonomous threat validation and cyber defense engineering. Founded and led by experienced red teamers who know that testing alone does not deliver better security, Cymulate goes beyond threat validation to build threat- and exposure-informed cyber defenses. For more information, visit www.cymulate.com.