INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-04-21 06:39 UTC

Frontier AI Compresses Patch-to-Exploit Windows: Security Teams Must Shift from Vulnerability Management to Exposure Prioritization

SECURITY ANALYSIS | **HIGH** | CVSS 7.5

SCC Item ID	SCC-STY-2026-0070
Type	Security Analysis
Severity	HIGH
CVSS Base Score	7.5
Affected Products	Enterprise Security Programs broadly; referenced platforms include CrowdStrike Falcon, Anthropic Claude Mythos, OpenAI GPT-5.4-Cyber (cybersecurity-specialized models)
Discovery Source	Rss:T1 Threatintel

Executive Summary

CrowdStrike's 2026 Global Threat Report documents a vendor-reported 89% year-over-year increase in AI-enabled adversary operations (confidence: medium, not independently audited) and a fastest-observed lateral movement breakout time of 27 seconds, figures that signal frontier AI is functionally eliminating the time window between vulnerability disclosure and active exploitation. The traditional scan-triage-patch cycle, built on the assumption that security teams have days or weeks to remediate disclosed vulnerabilities, no longer reflects operational reality for high-speed attack chains. Concurrently, defensive AI models from Anthropic and OpenAI are maturing and integrating into enterprise platforms, creating a parallel AI arms dynamic where defenders can leverage frontier models for detection and response to match the operational tempo of AI-assisted attacks.

Technical Analysis

CrowdStrike's 2026 Global Threat Report introduces three figures that, taken together, constitute a structural challenge to traditional vulnerability management: a vendor-reported 89% year-over-year increase in AI-enabled adversary operations (confidence: medium, not independently audited), a 42% rise in zero-days exploited prior to public disclosure (vendor data, confidence: medium), and a fastest-observed lateral movement breakout time of 27 seconds (single observation from CrowdStrike operational telemetry; not a statistical median or percentile). The directional signal is consistent with industry reporting on defensive AI announcements, and the implications

are operationally significant regardless of the exact figures.

The conventional vulnerability management model assumes a remediation window: a vulnerability is disclosed, a CVSS score is assigned, and security teams triage against severity. That model breaks down in two ways under AI-accelerated attack conditions. First, the 42% pre-disclosure exploitation rate means that for a meaningful fraction of vulnerabilities, no public advisory or patch exists when exploitation begins. CVSS-based triage cannot operate on a CVE that has not yet been published. Second, the 27-second breakout time compresses post-exploitation response to a window that manual incident response cannot match, regardless of triage speed.

The weakness classes most frequently weaponized in AI-accelerated chains, per the report, are CWE-269 (improper privilege management), CWE-732 (incorrect permission assignment), and CWE-306 (missing authentication). These are not novel or exotic weaknesses. They are structural gaps in identity and access configuration that AI-assisted tooling can enumerate and chain rapidly. The MITRE ATT&CK techniques associated with this activity pattern, T1068 (exploitation for privilege escalation), T1078 (valid accounts), T1550 (use of alternate authentication material), T1021 (remote services), T1190 (exploit public-facing application), T1651 (cloud administration command), T1595 (active scanning), T1110 (brute force), and T1588.006 (obtain capabilities: vulnerabilities), describe an end-to-end chain from initial reconnaissance through lateral movement and privilege escalation that AI tooling can execute and adapt at machine speed.

On the defensive side, two developments shift the calculus. Anthropic has announced Claude Mythos as a security-specialized frontier model, and CrowdStrike has joined the Mythos partner program. OpenAI has expanded access to cybersecurity-specialized models in response to the Mythos announcement. CrowdStrike has stated plans to integrate both into Falcon's AI-assisted detection and response workflows. This creates a parallel AI arms dynamic: adversaries using frontier models to accelerate exploitation, defenders using frontier models to accelerate detection and response.

The strategic implication is a forced transition from backlog-driven vulnerability management to continuous exposure prioritization. CISA's Stakeholder-Specific Vulnerability Categorization (SSVC) framework and CVSS exploitability metrics provide the methodological foundation for this transition. SSVC asks not 'how severe is this vulnerability?' but 'is this vulnerability being actively exploited, and does it affect systems with mission-critical exposure?' That question can be answered without waiting for a public CVE, which matters when 42% of zero-days are exploited before disclosure. Exposure prioritization frameworks also weight exploitability signals, active exploitation, proof-of-concept availability, threat actor targeting history, over raw severity scores, a reordering that aligns defensive effort with actual attack probability rather than theoretical impact.

Action Checklist

1. Step 1: Assess exposure, audit your current vulnerability management program against the assumption that remediation windows exist; document what percentage of your high-priority vulnerabilities have active exploit code or pre-disclosure exploitation history, and whether your tooling can surface that signal
2. Step 2: Review controls, verify identity and access configurations against CWE-269 (improper privilege management), CWE-732 (incorrect permission assignment), and CWE-306 (missing authentication); these are the weakness classes most frequently chained in AI-accelerated attacks; confirm MFA coverage on all remote access paths, enforce least privilege on service and admin accounts, and audit cloud administration command permissions against T1651
3. Step 3: Update threat model, incorporate AI-accelerated adversary operations as a threat scenario in your register; model the 27-second breakout time against your current detection and containment SLAs to

identify where automated response is required to meet acceptable response windows

4. Step 4: Evaluate prioritization methodology, if your program scores vulnerabilities by CVSS severity alone, assess adopting CISA's SSVC framework (https://csrc.nist.gov/publications/detail/sp/800-216/final) or an exploitability-weighted scoring model that can operate on pre-disclosure signals; map current tooling gaps against this requirement

5. Step 5: Communicate findings, brief leadership on the remediation window compression problem with specific context: not 'AI is changing cybersecurity' but 'our current SLA assumptions assume a window that the threat report says no longer reliably exists; here is what we need to close that gap'

6. Step 6: Monitor developments, track CrowdStrike's Falcon AI integration updates, Anthropic Mythos partner program announcements, and CISA SSVC guidance revisions; the defensive AI landscape is moving faster than annual review cycles

IR / Forensic Enrichment

Triage Priority	URGENT
Escalation Criteria	Escalate immediately to CISO and legal counsel if vulnerability management MTTR audit (Step 1) reveals that any KEV-listed vulnerability affecting internet-facing or identity systems has exceeded your documented SLA, or if the automated response gap analysis (Step 3) identifies that lateral movement containment cannot be achieved within a timeframe consistent with your regulatory breach notification obligations (e.g., 72-hour GDPR window, state breach notification statutes).
Recovery Notes	Recovery in the context of this threat is programmatic rather than system-specific: verify that all remediation SLA policies have been formally updated in writing to reflect exploitability-weighted prioritization and that the updated policies have been approved and distributed, not merely drafted. Monitor your first 90 days under the new prioritization model by tracking whether KEV-listed and high-EPSS vulnerabilities are being actioned faster than CVSS-only findings were previously — this delta is your proof-of-improvement metric. Continue watching the CrowdStrike Falcon AI integration roadmap and CISA SSVC guidance for material changes that would require another program update cycle, and plan a formal program review no later than 6 months from the date leadership is briefed.

Forensic Artifacts	Vulnerability management program baseline export: timestamped CSV of all open CVSS ≥7.0 findings cross-referenced against CISA KEV catalog, with MTTR metrics for the past 90 days — establishes the pre-remediation-window-compression posture and is required evidence if a subsequent AI-accelerated breach triggers regulatory inquiry into whether the CrowdStrike 2026 GTR findings were acted upon | Identity and access inventory snapshot: privileged account list with last-used timestamps, Windows Security Event Log Event ID 4672 (Special Privileges Assigned to New Logon) export for 30 days, and cloud IAM credential report (AWS or Azure native) — documents the attack surface exposed to AI-accelerated lateral movement exploiting CWE-269/CWE-732/CWE-306 weakness classes | Detection-to-containment latency record: empirical measurement of actual detection, alerting, and containment times from the past 10 security incidents, compared against the 27-second lateral movement breakout time documented in the CrowdStrike 2026 GTR — quantifies the response gap that drives the automated response requirement | MITRE ATT&CK Navigator layer file: current detective control coverage mapped against the Initial Access through Lateral Movement tactic chain, with T1651 (Cloud Administration Command) explicitly evaluated — provides a visual artifact of control gaps specific to AI-accelerated attack patterns documented in the 2026 GTR | Leadership brief package with date-stamped delivery record: the written gap analysis, updated SLA policy draft, and meeting record or email delivery confirmation — creates an auditable chain of evidence that the organization was formally informed of the remediation window compression problem and took documented action in response

Per-Action IR Details

Step 1: Assess exposure — audit your current vulnerability management program against the assumption that remediation windows exist; document what percentage of your high-priority vulnerabilities have active exploit code or pre-disclosure exploitation history, and whether your tooling can surface that signal

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Establishing IR capability and ensuring tooling reflects current threat reality

Controls: NIST SI-2 (Flaw Remediation), NIST RA-3 (Risk Assessment), NIST SI-5 (Security Alerts, Advisories, and Directives), CIS 7.1 (Establish and Maintain a Vulnerability Management Process), CIS 7.2 (Establish and Maintain a Remediation Process)

Compensating: Export your vulnerability scanner output (OpenVAS, Tenable Nessus Essentials, or Qualys Community) to CSV and cross-reference CVE IDs against CISA KEV catalog (https://www.cisa.gov/known-exploited-vulnerabilities-catalog) using a simple Python script or grep pipeline: `grep -Ff your_cve_list.txt kev_catalog.csv`. For pre-disclosure exploitation signals, subscribe to CISA's free AIS feed and NVD's JSON data feed via curl/cron. A 2-person team can automate this daily with a 30-line bash script piping KEV matches to email or Slack webhook — no SIEM required.

Evidence: Before restructuring your program, capture a baseline snapshot: export your current vulnerability scanner's open findings filtered to CVSS ≥7.0 as a timestamped CSV; pull your ticketing system's mean-time-to-remediate (MTTR) metrics for the past 90 days for critical/high findings; document your current SLA policy (e.g., 'critical vulns patched within 30 days'); and record what percentage of your scanner's findings include EPSS scores or KEV correlation. This baseline is your pre-change evidence record and will be required if a post-incident review asks whether your program's assumptions were documented before an AI-accelerated breach occurred.

Step 2: Review controls — verify identity and access configurations against CWE-269 (privilege management), CWE-732 (permission assignment), and CWE-306 (missing authentication); these are the weakness classes most frequently chained in AI-accelerated attacks; confirm MFA coverage on all remote access paths, enforce least privilege on service and admin accounts, and audit cloud administration command permissions against T1651

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Ensuring preventive controls reduce incident impact, specifically hardening the identity surface targeted by AI-accelerated lateral movement

Controls: NIST AC-2 (Account Management), NIST AC-6 (Least Privilege), NIST IA-2 (Identification and Authentication — Organizational Users), NIST IR-4 (Incident Handling), CIS 5.4 (Restrict Administrator Privileges to Dedicated Administrator Accounts), CIS 6.3 (Require MFA for Externally-Exposed Applications), CIS 6.4 (Require MFA for Remote Network Access), CIS 6.5 (Require MFA for Administrative Access), MITRE ATT&CK T1651 (Cloud Administration Command)

Compensating: Audit Windows service accounts using: `Get-ADServiceAccount -Filter * | Select Name, Enabled, LastLogonDate` and flag accounts with interactive logon rights via `Get-ADUser -Filter {ServicePrincipalNames -ne '$null'} -Properties ServicePrincipalNames, LastLogonDate`. For CWE-306 (missing authentication) on remote access paths, run `netstat -ano | findstr LISTENING` on internet-facing hosts and cross-reference open ports against expected services. For T1651 cloud command auditing without a commercial CSPM, use ScoutSuite (free, open-source multi-cloud auditor) to enumerate over-privileged IAM roles. For MFA gap identification, query Azure AD sign-in logs or export AWS IAM credential report (`aws iam generate-credential-report`) — both are free native capabilities.

Evidence: Capture before remediating: export a full privileged account inventory including service accounts, admin accounts, and cloud IAM roles with their last-used timestamps; pull Windows Security Event Log Event ID 4672 (Special Privileges Assigned to New Logon) for the past 30 days to identify accounts operating with SeDebugPrivilege or SeTcbPrivilege; enumerate cloud management plane API calls in AWS CloudTrail or Azure Activity Log filtered to administrative actions in the past 30 days; and document all remote access paths (VPN, RDP, SSH, cloud console) and their current MFA enrollment status. Given the 27-second lateral movement breakout documented in the CrowdStrike 2026 GTR, any unprotected admin credential represents a near-zero-dwell-time compromise path.

Step 3: Update threat model — incorporate AI-accelerated adversary operations as a threat scenario in your register; model the 27-second breakout time against your current detection and containment SLAs to identify where automated response is required to meet acceptable response windows

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Establishing IR capability aligned to current threat tempo, including automated response where human reaction time is insufficient

Controls: NIST IR-4 (Incident Handling), NIST IR-8 (Incident Response Plan), NIST RA-3 (Risk Assessment), NIST SI-4 (System Monitoring), CIS 7.1 (Establish and Maintain a Vulnerability Management Process)

Compensating: Quantify your detection-to-containment gap without a SIEM by measuring the manual steps in your current playbook: time your team performing each action in a tabletop against a 27-second lateral movement clock. Use the MITRE ATT&CK Navigator (free, browser-based) to map your current detective controls against the Initial Access → Lateral Movement kill chain and visually identify uncovered techniques. For automated containment without enterprise EDR, configure Windows Firewall via PowerShell to trigger host isolation on detection: `New-NetFirewallRule -DisplayName 'IR-Isolate' -Direction Inbound -Action Block -Enabled False` — pre-staged and enabled via script on alert. Document which response actions exceed human reaction time and require automation; this gap list is your business case for tooling investment.

Evidence: Before updating the threat model, collect empirical data on your current SLA performance: pull the last 10 security incidents from your ticketing system and calculate actual detection-to-containment times; query your endpoint agent logs (CrowdStrike Falcon sensor telemetry if deployed, or Sysmon EventID 1/3 logs) for mean time between initial execution and first lateral movement attempt in past incidents; and document your current alerting pipeline latency (time from log generation to analyst notification). This empirical baseline makes the 27-second breakout figure from the CrowdStrike 2026 GTR concrete and comparable against your actual response capability rather than theoretical SLAs.

Step 4: Evaluate prioritization methodology — if your program scores vulnerabilities by CVSS severity alone, assess adopting CISA's SSVC framework or an exploitability-weighted scoring model that can operate on pre-disclosure signals; map current tooling gaps against this requirement

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Ensuring detection and triage methodology reflects the compressed exploitation timeline documented in the CrowdStrike 2026 Global Threat Report

Controls: NIST SI-2 (Flaw Remediation), NIST SI-5 (Security Alerts, Advisories, and Directives), NIST RA-3 (Risk Assessment), CIS 7.1 (Establish and Maintain a Vulnerability Management Process), CIS 7.2 (Establish and Maintain a Remediation Process), CIS 7.3 (Perform Automated Operating System Patch Management), CIS 7.4 (Perform Automated Application Patch Management)

Compensating: Implement EPSS (Exploit Prediction Scoring System) scoring alongside CVSS at no cost: FIRST publishes daily EPSS scores in CSV format at https://www.first.org/epss/data_stats — download and join to your scanner output by CVE ID using a simple Python pandas merge. Add CISA KEV membership as a binary override field (KEV = immediate regardless of CVSS). This two-variable model (EPSS probability + KEV presence) approximates SSVC's 'exploitation' decision point with free data and a spreadsheet. For pre-disclosure signals without a threat intelligence platform, monitor vendor security advisory RSS feeds (Microsoft MSRC, CISA alerts) and configure a free RSS-to-email bridge so novel disclosures reach your team before scanner updates.

Evidence: Capture your current prioritization methodology in writing before changing it — screenshot or export the scoring logic in your existing scanner policy, document which fields drive ticket priority in your ticketing system, and pull a sample of the last 30 high/critical findings showing CVSS score, days-to-assignment, and days-to-remediation. This documents the pre-SSVC baseline and will be essential evidence if a regulator or insurer asks whether your program was updated in response to documented threat intelligence (specifically the CrowdStrike 2026 GTR's findings on AI-accelerated exploitation timelines).

Step 5: Communicate findings — brief leadership on the remediation window compression problem with specific context: not 'AI is changing cybersecurity' but 'our current SLA assumptions assume a window that the threat report says no longer reliably exists; here is what we need to close that gap'

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: Communicating lessons learned and program improvement requirements to leadership, and updating policies to reflect current threat intelligence

Controls: NIST IR-6 (Incident Reporting), NIST IR-8 (Incident Response Plan), NIST IR-2 (Incident Response Training), NIST SI-5 (Security Alerts, Advisories, and Directives), CIS 7.1 (Establish and Maintain a Vulnerability Management Process)

Compensating: Structure the leadership brief as a gap-to-consequence narrative using only public data: cite the CrowdStrike 2026 Global Threat Report (publicly available) for the 89% increase in AI-enabled operations and 27-second breakout figure; pair it with your MTTR baseline from Step 1 to show the arithmetic gap (e.g., 'our SLA is 30 days for critical vulns; the fastest observed exploitation now occurs in minutes'); and present three specific control investments ranked by cost-to-impact ratio. A 2-person team can build this brief in a single slide using the NIST CSF current-state vs. target-state gap format — framing the ask as 'moving from Tier 1 to Tier 2 on the CSF Respond function' gives leadership a recognized framework reference without requiring them to understand technical details.

Evidence: Before the brief, compile: the written program baseline from Step 1 (MTTR metrics, SLA policy document, KEV coverage percentage); the gap analysis from Step 4 (CVSS-only vs. exploitability-weighted scoring delta); and the automated response gap list from Step 3 (specific response actions that exceed human reaction time). These three documents transform the brief from a narrative claim into an evidence-backed risk presentation, and they create an auditable record that leadership was formally informed of the AI-accelerated threat model — relevant if a subsequent incident triggers regulatory or cyber insurance scrutiny.

Step 6: Monitor developments — track CrowdStrike's Falcon AI integration updates, Anthropic Mythos partner program announcements, and CISA SSVC guidance revisions; the defensive AI landscape is moving faster than annual review cycles

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: Updating threat intelligence processes and detection capabilities based on lessons learned from evolving adversary tooling

Controls: NIST SI-5 (Security Alerts, Advisories, and Directives), NIST IR-8 (Incident Response Plan), NIST AU-6 (Audit Record Review, Analysis, and Reporting), CIS 7.1 (Establish and Maintain a Vulnerability Management Process)

Compensating: Establish a free standing intelligence monitoring pipeline: subscribe to CISA Alerts RSS (https://www.cisa.gov/news-events/cybersecurity-advisories), CrowdStrike's Adversary Intelligence blog RSS, and MITRE ATT&CK changelog notifications via GitHub watch on the mitre-attack/attack-stac repository. Route all feeds into a free RSS aggregator (FreshRSS, self-hosted) or a Slack/Teams channel via webhook for team visibility. Schedule a 30-minute monthly review to assess whether new ATT&CK technique additions or CISA SSVC guidance changes require playbook updates — calendar-block this now rather than treating it as ad hoc. For CrowdStrike Falcon AI integration updates specifically, monitor the CrowdStrike release notes page and the Falcon community forums (free access with any Falcon license tier).

Evidence: Before establishing this monitoring pipeline, document the current state of your threat intelligence intake: list every external feed or subscription your team currently receives, when each was last reviewed, and whether any resulted in a program change in the past 12 months. This audit will likely reveal that annual review cycles are already inadequate for the current threat pace — the CrowdStrike 2026 GTR's 89% year-over-year increase in AI-enabled operations means a program reviewed in early 2025 may be materially out of date by mid-2025. The audit creates a documented basis for increasing review cadence and, if needed, justifying additional threat intelligence tooling investment.

Detection Guidance

Given the MITRE techniques associated with AI-accelerated attack chains in the report, detection should focus on behavioral anomalies rather than static signatures. Key areas to instrument:

Identity and access: Alert on T1078 (valid account abuse) patterns, logins at unusual hours, impossible travel, service account interactive logons, and token reuse consistent with T1550 (alternate authentication material: https://attack.mitre.org/techniques/T1550/). AI-assisted credential attacks (T1110: https://attack.mitre.org/techniques/T1110/) will generate authentication noise; tune alerting thresholds on failed login bursts against both on-premises and cloud targets.

Reconnaissance and initial access: T1595 (active scanning: https://attack.mitre.org/techniques/T1595/) and T1190 (exploit public-facing application: https://attack.mitre.org/techniques/T1190/) activity should be correlated against your external attack surface inventory. If you are not maintaining a continuous external exposure map, gaps here will be invisible. Review WAF and edge telemetry for scanning patterns that precede exploitation attempts.

Privilege escalation and lateral movement: T1068 (exploitation for privilege escalation: https://attack.mitre.org/techniques/T1068/) attempts should surface in EDR process telemetry as unexpected parent-child process relationships or privilege level transitions. T1021 (remote services: https://attack.mitre.org/techniques/T1021/) lateral movement at machine speed, the 27-second breakout window, will appear as rapid sequential authentication events across multiple hosts; ensure your SIEM correlation rules can fire on this pattern within seconds, not minutes.

Cloud and administration: T1651 (cloud administration command: https://attack.mitre.org/techniques/T1651/) abuse requires audit logging of cloud management plane actions; verify CloudTrail, Azure Activity Log, or equivalent is ingested and alerting on anomalous administrative API calls, particularly privilege escalation or data access actions from unfamiliar principals.

Weakness-class auditing: Run configuration audits specifically targeting CWE-269 (improper privilege management: https://cwe.mitre.org/data/definitions/269.html), CWE-732 (incorrect permission assignment: https://cwe.mitre.org/data/definitions/732.html), and CWE-306 (missing authentication: https://cwe.mitre.org/data/definitions/306.html). These are not detection rules, they are posture gaps that generate no alert when exploited if the underlying misconfiguration is present. Prioritize internet-facing systems and identity infrastructure.

Indicators of Compromise

Type	Value	Context	Confidence
TOOL	`Pending — refer to CrowdStrike 2026 Global Threat Report for published indicators`	The report documents AI-enabled attack chains leveraging T1068, T1078, T1550, T1021, and T1651 techniques; specific tooling, payload hashes, C2 infrastructure, or malware families associated with observed campaigns are not enumerated in the source material provided	LOW

Framework Mappings

MITRE-ATTACK

- **T1550** — Use Alternate Authentication Material
- **T1190** — Exploit Public-Facing Application
- **T1021** — Remote Services
- **T1068** — Exploitation for Privilege Escalation
- **T1078** — Valid Accounts
- **T1588.006** — Vulnerabilities
- **T1110** — Brute Force
- **T1651** — Cloud Administration Command
- **T1595** — Active Scanning

NIST-800-53R5

- **CA-8** — Penetration Testing
- **RA-5** — Vulnerability Monitoring and Scanning
- **SC-7** — Boundary Protection
- **SI-2** — Flaw Remediation
- **SI-7** — Software, Firmware, and Information Integrity
- **AC-17** — Remote Access
- **AC-3** — Access Enforcement
- **CM-7** — Least Functionality
- **IA-2** — Identification and Authentication (Organizational Users)
- **AC-6** — Least Privilege
- **AC-2** — Account Management
- **IA-5** — Authenticator Management
- **AC-7** — Unsuccessful Logon Attempts
- **CA-7** — Continuous Monitoring

- **SI-4** — System Monitoring
- **IR-5** — Incident Monitoring

OWASP-TOP10-2021

- **A01:2021** — Broken Access Control
- **A07:2021** — Identification and Authentication Failures

CIS-V8

- **5.4** — Restrict Administrator Privileges to Dedicated Administrator Accounts
- **6.8** — Define and Maintain Role-Based Access Control
- **6.3** — Require MFA for Externally-Exposed Applications
- **3.3** — Configure Data Access Control Lists
- **7.3** — Perform Automated Operating System Patch Management
- **7.4** — Perform Automated Application Patch Management
- **8.2** — Collect Audit Logs

ISO-27001-2022

- **A.8.8** — Management of technical vulnerabilities
- **A.5.21** — Managing information security in the ICT supply chain

SOC2-TSC

- **CC9.2** — Manages risks associated with vendors and business partners

HIPAA-SECURITY

- **164.312(e)(1)** — Transmission Security

NIST-CSF-2

- **DE.CM-01** — Networks and network services are monitored
- **DE.AE-08** — Incidents are declared when adverse events meet the defined incident criteria

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1550	Use Alternate Authentication Material	Defense-Evasion
T1190	Exploit Public-Facing Application	Initial-Access
T1021	Remote Services	Lateral-Movement
T1068	Exploitation for Privilege Escalation	Privilege-Escalation
T1078	Valid Accounts	Defense-Evasion
T1588.006	Vulnerabilities	Resource-Development
T1110	Brute Force	Credential-Access

Technique ID	Technique Name	Tactic
T1651	Cloud Administration Command	Execution
T1595	Active Scanning	Reconnaissance

Sources

Source	URL	Tier
Blog	https://www.crowdstrike.com/en-us/blog/frontier-ai-collapses-exploi...	T3
Frontier AI for Defenders: CrowdStrike and OpenAI TAC	https://www.crowdstrike.com/en-us/blog/frontier-ai-for-defenders-cr...	T3
In the Wake of Anthropic's Mythos, OpenAI Has a New Cybersecurity ...	https://www.wired.com/story/in-the-wake-of-anthropics-mythos-openai...	T2
OpenAI Widens Access to Cybersecurity Model After Anthropic's ...	https://www.securityweek.com/openai-widens-access-to-cybersecurity-...	T3
Anthropic Claude Mythos Preview - CrowdStrike	https://www.crowdstrike.com/en-us/blog/crowdstrike-founding-member-...	T3

Generated 2026-04-21 06:39 UTC by TJS Security Command Center