INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-04-28 13:43 UTC

Frontier AI Shrinks the Exploit Window to Near-Zero: Security Teams Must Abandon Backlog-Based Patching

SECURITY ANALYSIS | HIGH | CVSS 7.5

SCC Item ID	SCC-STY-2026-0090
Type	Security Analysis
Severity	HIGH
CVSS Base Score	7.5
Affected Products	Organizations using traditional severity-score-based vulnerability management programs; context references CrowdStrike Falcon Platform and frontier AI models (Anthropic Claude Mythos, OpenAI GPT-5.4-Cyber) as both threat-enabling and defensive tools
Discovery Source	Rss:T1 Threatintel

Executive Summary

Frontier AI models are collapsing the time between vulnerability disclosure and active exploitation to near-zero, challenging traditional severity-score-based patching programs. According to CrowdStrike's 2026 Global Threat Report, AI-enabled adversary attacks increased 89% year-over-year, pre-disclosure zero-day exploitation increased 42% year-over-year, and the fastest observed lateral movement breakout time was 27 seconds - figures that invalidate any patching workflow measured in days or weeks. The reported breach of Anthropic's Claude Mythos model compounds the signal: AI systems are simultaneously accelerating attacker capabilities and becoming targets themselves, forcing security leaders to rethink both their vulnerability management architecture and their AI supply chain risk posture.

Technical Analysis

The core thesis of CrowdStrike's 2026 Global Threat Report is structural, not episodic: frontier AI (large language models with advanced reasoning capabilities deployed in offensive tooling) has broken the defender assumption that a meaningful window exists between vulnerability disclosure and weaponization. That assumption underlies every CVSS-score-prioritized patching backlog in enterprise security programs today.

The reported metrics frame the severity of the shift. An 89% year-over-year increase in AI-enabled adversary attacks suggests the adoption curve for offensive AI tooling among threat actors is not incremental, it is accelerating. The 42% increase in zero-days exploited before public disclosure is particularly consequential: it means CVE publication, the trigger event for most patching workflows, is arriving after exploitation has already

begun in a growing share of cases. A 27-second lateral movement breakout time, the fastest CrowdStrike observed, is not a benchmark to defend against in isolation; it is an indicator of what automated, AI-assisted post-exploitation looks like when it operates with minimal operational overhead.

The MITRE ATT&CK techniques cited in this story map coherently to an AI-accelerated exploitation pattern. T1588.006 (Obtain Capabilities: Vulnerabilities) reflects adversary use of AI to rapidly scan, score, and prioritize vulnerabilities for weaponization. T1190 (Exploit Public-Facing Application) and T1212 (Exploitation for Credential Access) represent the initial access and credential harvesting phases that AI tooling can automate at scale. T1110 (Brute Force) reflects AI-assisted credential enumeration during initial access attempts. T1068 (Exploitation for Privilege Escalation), T1548 (Abuse Elevation Control Mechanism), and T1078 (Valid Accounts) reflect the post-exploitation chain that executes within seconds once initial access is achieved. T1072 (Software Deployment Tools) and T1210 (Exploitation of Remote Services) round out a lateral movement picture consistent with the 27-second breakout figure.

The CWEs mapped to this story, CWE-269 (Improper Privilege Management), CWE-732 (Incorrect Permission Assignment), CWE-862 (Missing Authorization), and CWE-306 (Missing Authentication for Critical Function), are not tied to a discrete vulnerability. They represent the authorization and access control weakness classes that become catastrophically exploitable when discovery-to-weaponization timelines compress to near-zero. Organizations with sprawling permission models, legacy systems carrying unreviewed access controls, and patching backlogs measured in weeks are the structural exposure AI-enabled adversaries are exploiting.

The CrowdStrike and OpenAI Threat Advisory Council partnership, along with CrowdStrike's position as a founding member of Anthropic's Claude Mythos program, signals that leading vendors are treating frontier AI as a defensive necessity, not an optional capability layer. Reported security concerns regarding the Mythos model, covered by Forbes and others, introduce a related risk dimension: if the AI models being integrated into security operations platforms are themselves compromised or face security incidents, the defensive tooling built on them carries inherited risk. Security teams evaluating AI-augmented detection and response platforms should include the security posture of the underlying AI system as part of their vendor risk assessment.

The strategic implication is clear: patching velocity must be decoupled from CVSS score queues and recoupled to real-time threat signal, specifically, exploitation likelihood and active adversary targeting, not theoretical severity. CISA's Known Exploited Vulnerabilities catalog and EPSS scoring offer more operationally relevant prioritization signals than CVSS base scores alone. Organizations that have not yet made this shift are operating a vulnerability management program calibrated to a threat tempo that no longer exists.

Action Checklist

1. Step 1: Assess exposure, audit whether your vulnerability management program prioritizes patching based primarily on CVSS base scores; if so, that workflow is the specific structural gap this story identifies
2. Step 2: Review controls, verify that your patching prioritization incorporates CISA KEV status and EPSS scores alongside CVSS; confirm EDR coverage and lateral movement detection are tuned for sub-minute breakout scenarios; review authentication controls against CWE-269, CWE-732, CWE-862, and CWE-306 weakness classes across identity and access management systems
3. Step 3: Update threat model, add AI-enabled adversary capability (automated vulnerability discovery, near-zero weaponization timelines, AI-assisted lateral movement) as a named threat class in your threat register; map against T1588.006, T1190, T1068, and T1072 for detection gap analysis
4. Step 4: Communicate findings, brief leadership that pre-disclosure zero-day exploitation is increasing and patching after CVE publication alone may not be a viable primary defense; frame the business risk as

a process design problem, not a staffing or tooling shortfall alone

5. Step 5: Monitor developments, track CrowdStrike's 2026 Global Threat Report follow-on publications, CISA KEV updates, and vendor communications regarding AI platform security incidents, particularly if your organization uses AI platforms with Anthropic model dependencies

6. Step 6: Review regulatory timelines, verify patching timelines comply with NIS2, CISA binding operational directives, or other applicable regulatory frameworks

IR / Forensic Enrichment

Triage Priority	URGENT
Escalation Criteria	Escalate immediately to CISO and legal counsel if: (1) any currently-open vulnerability in your backlog appears in a new CISA KEV entry within 24 hours of disclosure (indicating AI-accelerated weaponization has already occurred), (2) EDR or network logs show lateral movement events with sub-60-second dwell time consistent with the 27-second breakout documented in the CrowdStrike 2026 report, or (3) your organization has active API dependencies on Anthropic Claude Mythos and a confirmed breach disclosure is published — the latter triggers vendor notification obligations and potential supply chain incident classification under NIST IR-4 (Incident Handling) and applicable breach notification regulations if AI-processed data includes PII or PHI.
Recovery Notes	Recovery for a structural program failure of this type is process-level, not system-level: the primary recovery action is operationalizing a new prioritization workflow that gates patching decisions on KEV status and EPSS score before CVSS base score, and validating that EDR lateral movement detection rules are tuned to alert within 30 seconds of suspicious process-to-network chaining. Monitor the first 90 days of the new workflow for false-positive rate on EPSS-driven escalations and adjust the EPSS threshold (typically 0.1 or above is a reasonable starting gate) based on your team's triage capacity. If the alleged Anthropic Claude Mythos breach is confirmed during this period, conduct a focused supply chain review of all AI platform dependencies before returning any AI-assisted security tooling to production use.
Forensic Artifacts	CISA KEV JSON feed snapshot dated at time of program audit — provides timestamped proof of which vulnerabilities were known-exploited at the moment your backlog was evaluated, establishing whether any exploited CVEs were deprioritized due to score-based ordering | Windows Security Event Log Event ID 4688 (Process Creation) with full command-line logging enabled, filtered for lateral movement tool execution (psexec.exe, wmic.exe, net.exe, Invoke-Command) within 60-second windows of initial access events — directly evidences AI-assisted 27-second breakout behavior if it has already occurred in your environment | Network proxy or firewall flow logs showing outbound connections to known AI model provider API endpoints (api.anthropic.com, api.openai.com) with anomalous data transfer volumes — relevant if adversaries are using AI-as-a-Service for automated vulnerability weaponization originating from or targeting your network | EDR process tree captures (CrowdStrike Falcon process graph exports, or Sysmon Event ID 1 chains) showing parent-child relationships for any exploitation attempts against externally-exposed applications — maps directly to T1190 (Exploit Public-Facing Application) and T1068 (Exploitation for Privilege Escalation) technique artifact patterns | IAM audit logs from Active Directory, Entra ID, or Okta showing privilege escalation events (AD Event ID 4672 — Special Privileges Assigned to New Logon) time-correlated with network access events — the combination evidences automated post-exploitation privilege chaining consistent with AI-assisted lateral movement documented in the CrowdStrike 2026 report

Per-Action IR Details

Step 1: Assess exposure — audit whether your vulnerability management program prioritizes patching based primarily on CVSS base scores; if so, that workflow is the specific structural gap this story identifies

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Establishing IR Capability and Identifying Gaps in Current Controls

Controls: NIST SI-2 (Flaw Remediation) — requires testing and prioritizing software updates, not merely tracking scores, NIST RA-3 (Risk Assessment) — requires risk assessments that incorporate threat likelihood, not only CVSS base severity, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — process must be reviewed and updated annually or when significant threats emerge, CIS 7.2 (Establish and Maintain a Remediation Process) — explicitly requires risk-based remediation strategy, not score-based backlog ordering

Compensating: Export your current vuln backlog to CSV and run a cross-reference script against the live CISA KEV catalog (downloadable as JSON from cisa.gov/known-exploited-vulnerabilities-catalog) using a simple Python join on CVE ID. Any open finding matching a KEV entry that is CVSS-prioritized below critical in your queue is direct evidence of the structural gap. No SIEM required — this is a spreadsheet audit a single analyst can complete in under two hours.

Evidence: Before auditing the program, capture a timestamped export of your current vulnerability scanner results (Nessus, OpenVAS, or equivalent) sorted by CVSS base score descending — this becomes your baseline evidence of the score-based ordering. Also preserve the current KEV catalog snapshot (JSON, dated) and any EPSS score exports from FIRST.org for your open findings. These artifacts document the pre-remediation program state and are essential if you later need to demonstrate due diligence to auditors or regulators following an AI-accelerated exploitation event.

Step 2: Review controls — verify that your patching prioritization incorporates CISA KEV status and EPSS scores alongside CVSS; confirm EDR coverage and lateral movement detection are tuned for sub-minute breakout scenarios; review authentication controls against CWE-269, CWE-732, CWE-862, and CWE-306 weakness classes across identity and access management systems

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Tools, Techniques, and Controls Needed Before an Incident Occurs

Controls: NIST SI-4 (System Monitoring) — monitoring must be capable of detecting anomalous lateral movement at sub-minute timescales consistent with the 27-second breakout documented in the 2026 CrowdStrike Global Threat Report, NIST AC-6 (Least Privilege) — directly addresses CWE-269 (Improper Privilege Management) and CWE-732 (Incorrect Permission Assignment) by enforcing minimum necessary access, NIST AC-3 (Access Enforcement) — addresses CWE-862 (Missing Authorization) and CWE-306 (Missing Authentication for Critical Function) by enforcing authorization decisions at system boundaries, NIST IR-4 (Incident Handling) — requires preparation for the full range of incident types including those enabled by AI-accelerated adversary tooling, CIS 6.3 (Require MFA for Externally-Exposed Applications) — directly mitigates CWE-306 (Missing Authentication for Critical Function) on internet-facing surfaces, CIS 6.5 (Require MFA for Administrative Access) — addresses CWE-269 and CWE-862 by adding an authentication layer before privilege escalation can be completed, CIS 7.3 (Perform Automated Operating System Patch Management) — automated patching is prerequisite to any near-real-time remediation cadence required by AI-collapsed exploit windows

Compensating: For lateral movement detection without enterprise EDR: deploy Sysmon with SwiftOnSecurity's config (github.com/SwiftOnSecurity/sysmon-config) and enable Event ID 3 (Network Connection) and Event ID 1 (Process Creation) — alert on any process making lateral SMB or WinRM connections within 60 seconds of initial execution. For CWE-306 and CWE-862 coverage without a SIEM: run PowerShell query 'Get-ADUser -Filter * -Properties PasswordNeverExpires,LastLogonDate | Where {$_.PasswordNeverExpires -eq $true}' to surface accounts that could be abused without authentication enforcement. Use osquery scheduled query 'SELECT * FROM logged_in_users' with a 30-second interval to detect rapid session chaining consistent with 27-second breakout behavior.

Evidence: Capture EDR telemetry configuration exports (CrowdStrike Falcon sensor policy exports, or Sysmon XML config with hash) dated before this review — these prove your detection tuning baseline. Pull Windows Security Event Log Event ID 4624 (Logon) and 4625 (Failed Logon) for the past 30 days and filter on Logon Type 3 (network) and

Type 10 (RemoteInteractive) to identify accounts with authentication patterns consistent with CWE-306/CWE-862 exploitation paths. Export IAM permission reports from all identity providers (AD, Entra ID, Okta) — flag any accounts with AdminCount=1 that lack MFA enrollment, as these represent the highest-value targets for AI-assisted privilege escalation.

Step 3: Update threat model — add AI-enabled adversary capability (automated vulnerability discovery, near-zero weaponization timelines, AI-assisted lateral movement) as a named threat class in your threat register; map against T1588.006, T1190, T1068, and T1072 for detection gap analysis

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: Understanding Threat Actors, TTPs, and Building Detection Capability

Controls: NIST RA-3 (Risk Assessment) — threat register updates are a direct output of risk assessment and must reflect current threat intelligence including AI-enabled adversary capabilities documented in the 2026 CrowdStrike Global Threat Report, NIST SI-5 (Security Alerts, Advisories, and Directives) — organizations must receive and act on external threat intelligence; the 89% year-over-year increase in AI-enabled attacks constitutes an advisory-level signal requiring threat model update, NIST IR-5 (Incident Monitoring) — tracking and documenting incidents requires threat classification taxonomy that includes AI-enabled attack classes, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — vulnerability management process must incorporate threat intelligence about how vulnerabilities are being exploited, including AI-accelerated weaponization timelines

Compensating: Map your current detection rules against T1588.006 (Obtain Capabilities: Vulnerabilities — AI-assisted discovery), T1190 (Exploit Public-Facing Application), T1068 (Exploitation for Privilege Escalation), and T1072 (Software Deployment Tools) using the free MITRE ATT&CK Navigator (attack.mitre.org/resources/attack-navigator/). Export your Sigma rule inventory and filter for rules covering these four technique IDs — any gap is a detection hole that AI-enabled adversaries can exploit at near-zero weaponization timelines. For T1072 specifically, review any software deployment or RMM tools in your environment (PDQ Deploy, Ansible, SCCM) against the 2026 CrowdStrike report's documentation of adversary abuse of trusted deployment channels.

Evidence: Before updating the threat model, snapshot your current threat register and detection rule coverage report — this documents the pre-update state for audit and post-incident review purposes. Collect SIEM or Sysmon alert data for the past 90 days and filter for any events tagged to T1190 or T1068 to identify whether AI-enabled exploitation attempts have already occurred and been misclassified or missed. Pull network flow logs (NetFlow, Zeek connection.log, or Windows Firewall logs) and filter for repeated rapid-succession connection attempts to externally-exposed services — a pattern consistent with AI-automated vulnerability scanning described in the CrowdStrike 2026 report's T1588.006 abuse documentation.

Step 4: Communicate findings — brief leadership that the 42% increase in pre-disclosure zero-day exploitation means patching after CVE publication is no longer a viable primary defense; frame the business risk as a process design problem, not a staffing or tooling shortfall alone

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: Lessons Learned, Leadership Communication, and Program Improvement

Controls: NIST IR-6 (Incident Reporting) — requires personnel to report findings to organizational incident response capability and leadership; proactive communication about structural program gaps is within scope of IR-6, NIST IR-8 (Incident Response Plan) — the IR plan must be updated to reflect that CVE-publication-triggered patching is no longer sufficient as a primary defense given AI-collapsed weaponization timelines; leadership briefing is prerequisite to plan revision, NIST PM-9 (Risk Management Strategy) — leadership must understand changes to the risk environment; the 42% increase in pre-disclosure zero-day exploitation documented by CrowdStrike constitutes a material change requiring executive risk communication, CIS 7.2 (Establish and Maintain a Remediation Process) — process redesign from score-based to risk-based (KEV + EPSS + threat intelligence) requires leadership authorization and resource commitment that must be explicitly requested in the briefing

Compensating: Produce a one-page briefing document with three data points sourced directly from the CrowdStrike 2026 Global Threat Report: (1) 89% YoY increase in AI-enabled adversary attacks, (2) 42% increase in zero-days exploited before public disclosure, (3) 27-second fastest-observed lateral movement breakout. Frame each statistic against your organization's current mean-time-to-patch metric (pull this from your vuln scanner SLA reports). The gap between your patching timeline and these adversary timelines is the business risk figure — no additional tooling required to compute it.

Evidence: Before the leadership briefing, preserve the specific pages or sections of the CrowdStrike 2026 Global Threat Report that contain the 89%, 42%, and 27-second figures — these are your cited sources and must be retained for audit trail purposes. Also collect your organization's current documented patch SLA policy and any metrics showing actual mean-time-to-patch for the past two quarters. If EPSS data shows any currently-open findings with exploitation probability above 0.5 that have been sitting in the backlog behind lower-CVSS items, those specific findings are the most compelling concrete evidence that the process design problem exists today, not theoretically.

Step 5: Monitor developments — track CrowdStrike's 2026 Global Threat Report follow-on publications, CISA KEV updates, and any further disclosure regarding the alleged Anthropic Claude Mythos breach, particularly if your organization uses AI platforms with Anthropic model dependencies

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: Integrating Threat Intelligence into Ongoing Monitoring and Program Improvement

Controls: NIST SI-5 (Security Alerts, Advisories, and Directives) — requires ongoing receipt and action on external security advisories; CISA KEV updates and CrowdStrike threat report follow-on publications are authoritative external sources under this control, NIST IR-5 (Incident Monitoring) — tracking and documenting incidents requires continuous monitoring of threat intelligence feeds relevant to the organization's threat landscape, including AI platform supply chain risks, NIST SA-9 (External System Services) — organizations using Anthropic model APIs or AI platforms with Anthropic dependencies must monitor vendor security posture; any confirmed breach of Claude Mythos infrastructure triggers supply chain risk review under SA-9, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — vulnerability management process must incorporate continuous threat intelligence monitoring, including AI platform security advisories, not only NVD/CVE feeds

Compensating: Set up a free RSS or API-based alert pipeline for CISA KEV updates using the official KEV JSON feed (cisa.gov/sites/default/files/feeds/known_exploited_vulnerabilities.json) and poll it daily via a simple cron job or PowerShell scheduled task that diffs new entries against your asset inventory. For Anthropic Claude Mythos breach monitoring: subscribe to Anthropic's security advisories page and create a Google Alert for 'Anthropic security' and 'Claude Mythos breach' — if your organization has API contracts with Anthropic, review the vendor security notification clauses in those agreements now, before a breach is confirmed. For CrowdStrike follow-on publications, subscribe to the CrowdStrike Adversary Intelligence blog RSS feed as a no-cost continuous monitoring mechanism.

Evidence: If your organization uses AI platforms with Anthropic model dependencies, immediately inventory all API keys, OAuth tokens, and service account credentials used to authenticate to Anthropic endpoints — export these records with creation dates and last-used timestamps from your secrets manager or credential vault. Collect network proxy or firewall logs showing outbound connections to Anthropic API endpoints (api.anthropic.com) for the past 90 days to establish baseline traffic volume and data transfer sizes — anomalous outbound data volumes to AI provider endpoints could indicate data exfiltration via AI API abuse if a supply chain compromise is later confirmed. Preserve any AI-generated outputs or completions stored by your organization that were produced by Claude Mythos-lineage models, as these may become relevant if the alleged breach involved model poisoning or output manipulation rather than infrastructure compromise.

Detection Guidance

Traditional IOC-based detection is structurally insufficient against AI-accelerated exploitation; by the time indicators are published, the exploitation window has often closed or the attack has progressed. Behavioral and anomaly-based detection aligned to the observed MITRE techniques is the appropriate hunting frame.

For T1190 and T1212 (initial access and credential exploitation): Monitor authentication logs for rapid, sequential failed and successful login attempts against public-facing systems, particularly service accounts and privileged identities. Baseline normal authentication velocity and alert on deviations.

For T1068 and T1548 (privilege escalation): Hunt for privilege changes on accounts that have not historically held elevated permissions, particularly in short succession after initial authentication. Review SIEM rules for token manipulation and UAC bypass patterns.

For T1078 (valid accounts): Alert on first-use authentication from new devices or geographies for privileged accounts. Cross-correlate with vulnerability disclosure timelines; if a CVE drops and you see new authentication patterns within hours, treat that as a high-confidence exploitation signal.

For T1110 (brute force): Monitor for enumeration patterns targeting user directories or authentication services, particularly from external sources or newly compromised internal accounts, in the hours following vulnerability disclosure.

For T1072 (software deployment tools): Audit deployment and RMM tool invocation logs for lateral movement patterns, specifically, execution chains that originate from service accounts and traverse multiple hosts in rapid succession.

For T1210 (exploitation of remote services): Review network flow logs for unusual internal scanning or connection attempts to services not normally accessed by the originating host, particularly within the first 24 hours after a CVE disclosure.

For AI-enabled adversary tempo specifically: The 27-second breakout figure means detection and containment playbooks that require human-in-the-loop approval at each step will fail to contain active intrusions. Review whether automated containment actions (host isolation, credential suspension) are authorized in your EDR and SOAR platforms, and under what threshold conditions they fire.

For the Mythos security concern: If your organization uses AI platforms with Anthropic model dependencies or CrowdStrike Falcon with AI-augmented features, review vendor communications for any service advisories. Audit API key and credential access for AI platform integrations.

Indicators of Compromise

Type	Value	Context	Confidence
TOOL	`Pending — refer to CrowdStrike 2026 Global Threat Report and associated blog publications for any published indicators`	CrowdStrike's reporting on AI-enabled adversary attacks references observed TTPs and campaign behaviors; specific IOC values (hashes, IPs, domains) are not enumerated in the source material provided	LOW

Framework Mappings

MITRE-ATTACK

- **T1588.006** — Vulnerabilities
- **T1190** — Exploit Public-Facing Application
- **T1212** — Exploitation for Credential Access

- **T1078** — Valid Accounts
- **T1110** — Brute Force
- **T1548** — Abuse Elevation Control Mechanism
- **T1072** — Software Deployment Tools
- **T1210** — Exploitation of Remote Services
- **T1203** — Exploitation for Client Execution
- **T1068** — Exploitation for Privilege Escalation

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-2** — Account Management
- **AC-6** — Least Privilege
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **AC-7** — Unsuccessful Logon Attempts
- **CM-6** — Configuration Settings
- **SI-3** — Malicious Code Protection
- **SI-4** — System Monitoring
- **AC-3** — Access Enforcement
- **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
- **6.1** — Establish an Access Granting Process
- **3.3** — Configure Data Access Control Lists
- **6.3** — Require MFA for Externally-Exposed Applications
- **7.3** — Perform Automated Operating System Patch Management
- **7.4** — Perform Automated Application Patch Management

ISO-27001-2022

- **A.8.8** — Management of technical vulnerabilities

NIST-CSF-2

- **DE.AE-08** — Incidents are declared when adverse events meet the defined incident criteria

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1588.006	Vulnerabilities	Resource-Development
T1190	Exploit Public-Facing Application	Initial-Access
T1212	Exploitation for Credential Access	Credential-Access
T1078	Valid Accounts	Defense-Evasion
T1110	Brute Force	Credential-Access
T1548	Abuse Elevation Control Mechanism	Privilege-Escalation
T1072	Software Deployment Tools	Execution
T1210	Exploitation of Remote Services	Lateral-Movement
T1203	Exploitation for Client Execution	Execution
T1068	Exploitation for Privilege Escalation	Privilege-Escalation

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
Mythos Is a Wake-Up Call: Five Steps to Prepare for Frontier AI	https://www.crowdstrike.com/en-us/resources/crowdcasts/mythos-is-a-...	T3
Anthropic Claude Mythos Preview - CrowdStrike	https://www.crowdstrike.com/en-us/blog/crowdstrike-founding-member-...	T3
Alleged Claude Mythos Breach Raises Questions About AI Security	https://www.forbes.com/sites/timkeary/2026/04/23/alleged-claude-myt...	T3

Generated 2026-04-28 13:43 UTC by TJS Security Command Center