INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-05-01 14:03 UTC

Frontier AI Is Closing the Exploit Window: Traditional Vulnerability Management Can No Longer Keep Pace

SECURITY ANALYSIS | HIGH | CVSS 7.5

SCC Item ID	SCC-STY-2026-0102
Type	Security Analysis
Severity	HIGH
CVSS Base Score	7.5
Affected Products	General, organizations relying on periodic vulnerability scanning and CVSS-driven prioritization; CrowdStrike Falcon Platform referenced as detection/response context
Discovery Source	Rss:T1 Threatintel

Executive Summary

Frontier AI is compressing the time between vulnerability disclosure and active exploitation to near-real-time, fundamentally breaking the assumption that defenders have a meaningful patch window. According to CrowdStrike's 2026 Global Threat Report, AI-enabled adversary attacks increased 89% year-over-year, with fastest-observed lateral movement breakout time of 27 seconds - figures that render periodic scanning cycles and CVSS-driven prioritization structurally inadequate. For CISOs and boards, this is not a technology refresh question; it is an operating model question: security programs built on human-paced detection and response timelines are misaligned with the speed at which AI-augmented adversaries now operate.

Technical Analysis

The central argument in CrowdStrike's 2026 Global Threat Report, and the supporting blog series on frontier AI, is that AI is not merely automating existing attacker workflows, it is collapsing the temporal advantage defenders historically relied on. Two data points anchor the analysis: a 42% increase in zero-days exploited before public disclosure, and a 27-second lateral movement breakout time. Together, these figures describe an environment where the traditional vulnerability management lifecycle - scan, score by CVSS, prioritize, patch, verify - cannot close the gap between exposure and exploitation.

The MITRE ATT&CK techniques mapped to this threat pattern illustrate the full kill chain AI-augmented adversaries are accelerating. Initial access via public-facing application exploitation (T1190) or valid account abuse (T1078) is followed by rapid lateral movement (T1021) and privilege escalation (T1068). Credential

brute-forcing (T1110) and use of stolen credentials or web session tokens (T1550) reduce the time needed to establish persistence. Critically, impair defenses (T1562) activity, disabling logging, EDR agents, or alerting pipelines, is appearing earlier in the attack sequence, suggesting adversaries are using AI to identify and neutralize detection capabilities before executing their primary objective. Tool and vulnerability weaponization acquisition (T1588.006) and exploitation of remote services (T1210) round out a pattern where every phase of the attack lifecycle is being accelerated.

The mapped CWEs (CWE-284, CWE-287, CWE-250, CWE-269) - improper access control, improper authentication, excessive privilege, and improper privilege management - describe the structural weaknesses AI-assisted attackers are most effectively targeting. These are editorial mappings to the threat pattern described, not tied to a specific CVE. Over-privileged service accounts, incomplete MFA enforcement, and excessive standing access provide the footholds that fast-moving adversaries exploit before defenders can respond.

One significant sourcing limitation applies to this story: secondary sources reference product names 'Anthropic Claude Mythos' and 'OpenAI GPT-5.4-Cyber,' which cannot be independently verified as commercial product identifiers. These product-level claims should be independently verified against official Anthropic and OpenAI documentation before use in formal communications. The core statistical claims - 89% increase in AI-enabled attacks, 42% increase in pre-disclosure zero-day exploitation, 27-second breakout time - are attributed to CrowdStrike's 2026 Global Threat Report and should be verified against that primary document before use in formal communications. This item is not tied to a specific CVE; the qualitative_rating of 'high' reflects editorial assessment of the threat pattern described.

Action Checklist

1. Step 1: Assess exposure, audit whether your vulnerability management program operates on periodic scan cycles (weekly, monthly) rather than continuous monitoring; if so, quantify the gap between scan cadence and current mean time to exploit for your most common vulnerability classes
2. Step 2: Review controls, verify MFA enforcement across all privileged and remote access paths (T1078, T1110); audit EDR agent coverage for gaps and confirm that impair-defenses detections (T1562) are alerting correctly; review service account privilege levels against least-privilege standards (CWE-250, CWE-269)
3. Step 3: Update threat model, add AI-accelerated exploitation as a named threat pattern in your threat register; specifically document the assumption that a meaningful patch window no longer exists for internet-facing systems and high-value internal assets; map to T1190, T1068, T1021, T1562
4. Step 4: Communicate findings, brief leadership using the 27-second breakout time and 42% pre-disclosure zero-day figures as concrete anchors; frame the business question as: what is our detection-to-containment time, and does it beat 27 seconds for lateral movement scenarios
5. Step 5: Monitor developments and verify claims. Track CrowdStrike's 2026 Global Threat Report for release of full statistical methodology and attack campaign attribution. Prioritize independent corroboration of the AI-enabled attack metrics from CISA, NIST exploitation tracking, or peer vendor threat intelligence before using figures in board-level communications.

IR / Forensic Enrichment

Triage Priority	URGENT
Escalation Criteria	Escalate to CISO and executive leadership immediately if detection-to-containment time measurement (Step 4) reveals D2C exceeds 30 minutes for lateral movement scenarios, if EDR coverage gaps (Step 2) expose more than 10% of internet-facing or Tier-0 assets without agent coverage, or if CISA issues a KEV entry for a vulnerability present in your environment while your organization remains on a weekly-or-slower scan cycle — any of these conditions indicates structural inability to respond within the AI-accelerated exploitation window documented in the CrowdStrike 2026 GTR.
Recovery Notes	Recovery in the context of this threat story is programmatic rather than system-specific: after implementing continuous monitoring and tightening MFA and EDR coverage per Steps 1 and 2, verify recovery by re-running the detection-to-containment time calculation from Step 4 against a controlled tabletop or purple team exercise scripted to the 27-second lateral movement scenario, confirming D2C has measurably decreased. Monitor your vulnerability management SLA compliance weekly for 90 days post-implementation to confirm the shift from periodic to continuous scanning is operationally sustained and not reverting under workload pressure. Treat any future CISA KEV entry for software in your environment as a live test of the updated program: measure time from KEV publication to confirmed remediation or compensating control deployment and track this as a KPI to demonstrate improvement to leadership.
Forensic Artifacts	Windows Security Event Log — Event ID 4648 (Logon with explicit credentials) and 4624 Type 3 (Network logon) on domain controllers and Tier-0 servers: AI-accelerated lateral movement via T1021 will produce a statistically anomalous burst of these events within a compressed timeframe (seconds to low minutes) rather than the distributed pattern of human-speed lateral movement; baseline normal hourly rates before any incident to make the anomaly detectable | EDR process tree telemetry (CrowdStrike Falcon Process Timeline or equivalent) for processes spawned by internet-facing service accounts: T1190 exploitation followed by T1068 privilege escalation will produce characteristic parent-child process relationships (e.g., web server process spawning cmd.exe, powershell.exe, or mshta.exe) that represent the initial post-exploitation foothold before AI-assisted lateral movement begins at T1021 | CrowdStrike Falcon Incident Workbench 'Impair Defenses' detection alerts (T1562): AI-assisted exploit chains specifically target EDR agent disable/tamper as an early kill-chain step to eliminate telemetry before lateral movement; the presence or conspicuous absence of T1562 alerts during a suspected intrusion window is itself a forensic indicator | CISA KEV catalog delta log — a versioned daily snapshot of the KEV JSON feed cross-referenced against your asset inventory: when an exploitation event is investigated, the timestamp of KEV listing relative to your last successful scan of the affected asset quantifies exactly how large your exposure window was, providing concrete forensic evidence of the patch-window gap this threat story describes | Active Directory Kerberos TGS request logs (Event ID 4769 on domain controllers) filtered for service accounts with weak encryption types (RC4, 0x17): AI-accelerated Kerberoasting (T1558.003) as a precursor to T1078 valid-account abuse will produce high-volume TGS requests for service accounts in a compressed window; these logs must be captured before any credential rotation activity in Step 2 to preserve the pre-remediation forensic baseline

Per-Action IR Details

Step 1: Assess exposure — audit whether your vulnerability management program operates on periodic scan cycles (weekly, monthly) rather than continuous monitoring; if so, quantify the gap between scan cadence and current mean time to exploit for your most common vulnerability classes

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Establishing IR capability, tooling, and visibility baselines before adversary activity begins

Controls: NIST SI-4 (System Monitoring) — continuous monitoring requirement directly addresses the cadence gap this step quantifies, NIST RA-3 (Risk Assessment) — requires organizations to assess likelihood and impact, which must account for near-zero patch windows under AI-accelerated exploitation, NIST SI-2 (Flaw Remediation) — flaw remediation timelines must be reassessed when MTTE for common vulnerability classes now approaches or precedes vendor patch availability, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — explicitly requires documenting scan frequency and remediation SLAs; this step operationalizes a review of those documented commitments against current threat reality, CIS 7.2 (Establish and Maintain a Remediation Process) — risk-based remediation strategy must be recalibrated when 42% of zero-days are exploited pre-disclosure, making CVSS-score-at-disclosure an inadequate trigger

Compensating: For teams without a continuous scanning platform: deploy OpenVAS (Greenbone Community Edition) in authenticated scan mode on a 24-hour cron schedule targeting internet-facing assets and domain controllers; run `nmap -sV --script vulners -p 80,443,8080,8443,445,3389,22` nightly against your perimeter and pipe output to a diff script that alerts on new open ports or version changes. Use CISA's KEV catalog (downloaded via `curl https://www.cisa.gov/sites/default/files/feeds/known_exploited_vulnerabilities.json`) cross-referenced against your asset inventory in a spreadsheet to identify KEV-listed CVEs present in your environment — this approximates continuous prioritization without a SIEM.

Evidence: Before reconfiguring scan cadence, capture a point-in-time baseline: export your current vulnerability scanner's last full scan report with timestamps to establish the historical scan interval; pull CISA KEV JSON and record the `dateAdded` vs. `dueDate` delta for any CVEs matching your installed software to quantify how often exploitation precedes your scan window; document your scanner's authenticated vs. unauthenticated coverage ratio, as unauthenticated scans systematically miss privilege-escalation-class vulnerabilities (T1068) that AI-assisted exploit chains favor for rapid lateral movement.

Step 2: Review controls — verify MFA enforcement across all privileged and remote access paths (T1078, T1110); audit EDR agent coverage for gaps and confirm that impair-defenses detections (T1562) are alerting correctly; review service account privilege levels against least-privilege standards (CWE-250, CWE-269)

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Ensuring detection and defensive controls are operational and correctly tuned before an AI-accelerated intrusion attempt occurs

Controls: NIST AC-2 (Account Management) — service account enumeration and privilege review maps directly to CWE-269 (Improper Privilege Management); accounts with excessive rights are the primary escalation path in T1078 abuse, NIST IA-5 (Authenticator Management) — MFA enforcement verification across privileged and remote paths addresses T1110 (Brute Force) and T1078 (Valid Accounts), both of which are accelerated by AI-driven credential stuffing and password spraying at machine speed, NIST SI-4 (System Monitoring) — T1562 (Impair Defenses) detection alerting verification is a direct SI-4 implementation check; if EDR agent tampering goes undetected, the 27-second lateral movement window becomes irrelevant because containment telemetry disappears, NIST IR-4 (Incident Handling) — EDR coverage gap audit is preparation-phase IR-4 work: you cannot execute containment at 27-second breakout speed if agent gaps leave blind spots on lateral movement paths, CIS 6.3 (Require MFA for Externally-Exposed Applications) — direct control for the T1078/T1110 MFA verification task, CIS 6.5 (Require MFA for Administrative Access) — privileged path MFA audit; AI-enabled adversaries specifically target unprotected admin accounts because compromising one collapses the kill chain, CIS 5.4 (Restrict Administrator Privileges to Dedicated Administrator Accounts) — service account least-privilege review operationalizes this safeguard against CWE-250 and CWE-269

Compensating: MFA gap audit without enterprise IAM tooling: run `Get-ADUser -Filter * -Properties 'msDS-SupportedEncryptionTypes','ServicePrincipalName','PasswordLastSet' | Where-Object {$_.ServicePrincipalName -ne $null}` to enumerate service accounts, then cross-reference against your MFA enrollment list manually. For EDR coverage gaps without a fleet management console, deploy osquery with the CIS osquery pack and query `SELECT * FROM processes WHERE name = 'MsMpEng.exe' OR name = 'CSFalconService.exe'` across endpoints via scheduled task; missing results identify uncovered hosts. For T1562 alerting verification, deploy the Sigma rule `windows_defender_tamper` (SigmaHQ detection-rules repo) and test by

temporarily disabling real-time protection on a lab host to confirm alert fires.

Evidence: Before auditing controls, collect: Windows Security Event Log Event ID 4625 (failed logons) and 4648 (explicit credential use) from domain controllers for the trailing 30 days to establish T1110 baseline activity volume — AI-driven credential attacks will appear as statistically anomalous spikes against this baseline; CrowdStrike Falcon (or equivalent EDR) agent heartbeat logs showing last check-in time per host to identify stale/missing agents before the audit modifies any configuration; Active Directory replication metadata (`repadmin /showrepl`) to identify service accounts with `PasswordNeverExpires` and `DontRequirePreauth` flags set, as these are the highest-value targets for AI-accelerated Kerberoasting (T1558.003) chained into T1078.

Step 3: Update threat model — add AI-accelerated exploitation as a named threat pattern in your threat register; specifically document the assumption that a meaningful patch window no longer exists for internet-facing systems and high-value internal assets; map to T1190, T1068, T1021, T1562

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Updating threat models and IR plans to reflect current adversary capability, specifically the elimination of the assumed patch window documented in CrowdStrike's 2026 Global Threat Report

Controls: NIST RA-3 (Risk Assessment) — threat model update is a direct RA-3 activity; the specific update required here is adding 'AI-accelerated exploitation with near-zero patch window' as a named likelihood driver that elevates risk ratings for internet-facing assets independent of CVSS score, NIST RA-5 (Vulnerability Monitoring and Scanning) — the threat model must now reflect that RA-5 scan frequency is no longer sufficient as a primary risk reduction control; this updates the compensating control assumptions embedded in existing risk assessments, NIST IR-8 (Incident Response Plan) — IR plan must be updated to reflect T1190/T1068 exploitation scenarios where the window between disclosure and active exploitation is measured in hours, not weeks; playbooks built on 'patch within 30 days' SLAs require immediate revision, NIST SI-5 (Security Alerts, Advisories, and Directives) — AI-accelerated exploitation changes the operationalization of SI-5: advisories can no longer be queued for weekly triage; the threat model update must encode a near-real-time advisory response workflow, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — the vulnerability management process documentation must be revised to name AI-assisted exploitation as a process assumption that invalidates periodic-scan-only approaches

Compensating: For teams maintaining threat registers in spreadsheets rather than GRC platforms: add a column `AI_Accel_Applicable` (boolean) to your existing threat register and mark true for any threat scenario involving T1190, T1068, T1021, or T1562; add a second column `Patch_Window_Assumption_Days` and update internet-facing asset rows to 0 (no assumed window) per CrowdStrike 2026 GTR findings. For ATT&CK technique mapping without a commercial threat intelligence platform, use the MITRE ATT&CK Navigator (free, browser-based at attack.mitre.org/resources/attack-navigator/) to build a layer file for T1190, T1068, T1021.001, T1021.002, and T1562.001, then export as JSON to attach to your threat register entry.

Evidence: Before finalizing threat model updates, preserve the current state as a versioned baseline: export your existing threat register with timestamps and current risk ratings for all internet-facing asset threat scenarios; document current patch SLA commitments from your vulnerability management policy (the specific number of days); pull CISA KEV `dateAdded` vs. NVD `publishedDate` delta for the last 90 days of entries to empirically quantify how often exploitation precedes or immediately follows disclosure in your relevant software categories — this data substantiates the threat model assumption change with evidence rather than vendor claims alone.

Step 4: Communicate findings — brief leadership using the 27-second breakout time and 42% pre-disclosure zero-day figures as concrete anchors; frame the business question as: what is our detection-to-containment time, and does it beat 27 seconds for lateral movement scenarios

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: Lessons learned and capability gap reporting to leadership; this step applies post-incident communication discipline to a proactive capability assessment, using published threat data as a surrogate for observed incident data

Controls: NIST IR-6 (Incident Reporting) — while this is a proactive brief rather than an active incident report, IR-6's requirement to communicate incident-relevant information to appropriate organizational personnel directly applies;

leadership needs the 27-second breakout metric to make resourcing decisions before the incident occurs, NIST IR-8 (Incident Response Plan) — the brief should result in documented updates to the IR plan authorizing faster containment actions (network isolation, account lockout) at lower confidence thresholds, given that waiting for confirmation is structurally incompatible with 27-second lateral movement, NIST RA-3 (Risk Assessment) — the 42% pre-disclosure zero-day figure and 89% YoY increase in AI-enabled attacks are quantitative inputs to RA-3 likelihood determinations; the brief should frame these as risk assessment data points requiring plan updates, not just awareness, CIS 7.2 (Establish and Maintain a Remediation Process) — the business question framing ('does our detection-to-containment time beat 27 seconds') is a direct challenge to existing remediation SLA assumptions that CIS 7.2 requires be documented and reviewed

Compensating: To calculate actual detection-to-containment time without a SIEM: pull the five most recent security incidents from your ticketing system, calculate elapsed time between 'alert created' and 'host isolated/account locked' timestamps manually, and present the median as your current D2C time against the 27-second benchmark. For teams with no prior incident history, run a tabletop exercise specifically scripted around a 27-second lateral movement scenario (attacker compromises one host, moves to domain controller within 27 seconds) using the CISA Tabletop Exercise Packages (CTEPs) framework to generate an estimated response time for the brief.

Evidence: Before the leadership brief, gather your own organizational data to contextualize vendor figures: pull EDR or Windows Event Log mean time between alert generation (Event ID 1102 — audit log cleared, as a proxy for attacker anti-forensics activity, or EDR alert timestamps) and analyst acknowledgment from the last 90 days; document the current number of internet-facing systems without continuous monitoring coverage; if CrowdStrike Falcon is deployed, pull the 'Time to Detect' and 'Time to Investigate' metrics from the Falcon console's Incident Workbench for recent detections — these become your organization-specific counterpart to the 27-second GTR figure and are more persuasive to leadership than vendor benchmarks alone.

Step 5: Monitor developments — track CrowdStrike's 2026 Global Threat Report release for the full statistical methodology; watch for independent corroboration of the AI-enabled attack metrics from CISA, NIST NVD exploitation data, or peer vendor threat intelligence

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: Continuous improvement through integration of updated threat intelligence; specifically, monitoring for corroboration of AI-acceleration metrics maps to the lessons-learned and threat model refinement activities NIST 800-61r3 assigns to this phase

Controls: NIST SI-5 (Security Alerts, Advisories, and Directives) — SI-5 requires receiving and acting on security advisories from external organizations on an ongoing basis; this step operationalizes SI-5 for the specific case of AI-acceleration threat intelligence, naming CISA and NVD as the authoritative external sources to monitor, NIST IR-5 (Incident Monitoring) — tracking the evolution of AI-enabled attack metrics over time is an IR-5 activity; the threat register entry created in Step 3 must be updated as corroborating data arrives from CISA KEV, NVD exploitation flags, and peer vendor reports, NIST DE.AE-07 (Cyber threat intelligence integrated into adverse event analysis) — this NIST 800-61r3 CSF-aligned recommendation directly requires integrating up-to-date CTI into adverse event analysis; monitoring for corroboration of the 89% AI-attack increase and 27-second breakout metrics is how you keep that CTI current, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — the vulnerability management process must be updated when corroborating data changes the empirical basis for patch window assumptions; this monitoring step creates the trigger mechanism for that update

Compensating: For teams without a commercial threat intelligence feed: subscribe to CISA's free advisories via RSS (https://www.cisa.gov/cybersecurity-advisories — verify URL at time of use) and set a calendar reminder to cross-reference NVD's exploitation data (`https://nvd.nist.gov/vuln/search` with `hasKev=1` filter — verify URL at time of use) monthly; create a shared document tracking the following metrics over time: CISA KEV entries per month, median days-to-KEV-listing from NVD publish date, and any CISA/NSA joint advisory language referencing AI-assisted exploitation. Use the MITRE ATT&CK changelog RSS feed to track additions of AI-assisted technique sub-procedures as independent corroboration of the threat pattern.

Evidence: Establish a monitoring baseline before tracking begins: snapshot the current CISA KEV catalog entry count and the current NVD count of CVEs flagged with EPSS scores above 0.9 (indicating high exploitation probability); record the current version and publication date of the CrowdStrike 2026 Global Threat Report preview material being referenced; document which peer vendor threat reports (Mandiant M-Trends, Verizon DBIR, Microsoft MSTIC annual

report) are currently in your intelligence review cycle so that additions can be tracked as corroborating sources when AI-acceleration metrics appear in those independent publications.

Detection Guidance

Given the attack pattern described, detection focus should shift from indicator-matching to behavior sequencing at machine speed. Specific areas to instrument and hunt:

- Lateral movement velocity: Alert on authentication events or remote service connections (T1021) occurring within seconds of an initial access event or privilege escalation. A 27-second breakout time means dwell-time-based detection thresholds calibrated for minutes or hours will miss these sequences entirely.

- Defense impairment sequencing (T1562): Hunt for EDR agent stops, log forwarding interruptions, or firewall rule modifications occurring within the first minutes of a new session or elevated process. Early impairment activity is a strong behavioral indicator of AI-assisted or scripted attack execution.

- Privilege escalation followed by lateral movement (T1068 → T1021): Correlate local privilege escalation events with subsequent remote service connections, particularly to high-value targets such as domain controllers, backup infrastructure, or payment systems.

- Valid account abuse at speed (T1078, T1550): Monitor for credential reuse or session token replay across multiple systems in compressed time windows. AI-assisted credential stuffing (T1110) will produce authentication attempt volumes and velocities inconsistent with human behavior.

- Pre-disclosure zero-day exploitation (T1190): Because 42% of zero-days in this reporting period were exploited before public disclosure, signature-based detection on CVE identifiers will fail. Prioritize behavioral detection on post-exploitation activity rather than exploit signatures.

- Log sources to prioritize: authentication logs (Windows Security Event ID 4624, 4648, 4625), EDR process telemetry, network flow data for east-west lateral movement, and endpoint integrity monitoring for agent status changes.

Security operations teams should review detection rule coverage for the mapped ATT&CK techniques, particularly T1562 (impair defenses) and T1068 (exploitation for privilege escalation), and validate that behavioral analytics modules are active and tuned to catch AI-assisted attack sequences.

Indicators of Compromise

Type	Value	Context	Confidence
TOOL	`Pending — refer to CrowdStrike 2026 Global Threat Report for published indicators`	CrowdStrike's 2026 Global Threat Report is cited as the primary source for AI-enabled attack statistics; specific IOCs, tool hashes, and campaign indicators associated with documented AI-assisted intrusions may be published in that report or its accompanying intelligence releases	LOW

Framework Mappings

MITRE-ATTACK

- **T1210** — Exploitation of Remote Services
- **T1550** — Use Alternate Authentication Material
- **T1190** — Exploit Public-Facing Application
- **T1562** — Impair Defenses
- **T1078** — Valid Accounts
- **T1110** — Brute Force
- **T1021** — Remote Services
- **T1068** — Exploitation for Privilege Escalation
- **T1588.006** — Vulnerabilities

NIST-800-53R5

- **AC-6** — Least Privilege
- **SC-7** — Boundary Protection
- **SI-2** — Flaw Remediation
- **CA-8** — Penetration Testing
- **RA-5** — Vulnerability Monitoring and Scanning
- **SI-7** — Software, Firmware, and Information Integrity
- **AU-9** — Protection of Audit Information
- **CM-6** — Configuration Settings
- **SI-4** — System Monitoring
- **AC-2** — Account Management
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **AC-7** — Unsuccessful Logon Attempts
- **AC-17** — Remote Access
- **AC-3** — Access Enforcement
- **CM-7** — Least Functionality
- **IA-8** — Identification and Authentication (Non-Organizational Users)
- **IR-5** — Incident Monitoring

OWASP-TOP10-2021

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

CIS-V8

- **6.1** — Establish an Access Granting Process
- **6.2** — Establish an Access Revoking Process
- **6.3** — Require MFA for Externally-Exposed Applications
- **6.4** — Require MFA for Remote Network Access

- **6.5** — Require MFA for Administrative Access
- **5.4** — Restrict Administrator Privileges to Dedicated Administrator Accounts
- **6.8** — Define and Maintain Role-Based Access Control
- **7.3** — Perform Automated Operating System Patch Management
- **7.4** — Perform Automated Application Patch Management

SOC2-TSC

- **CC6.1** — The entity implements logical access security software, infrastructure, and architectures over protected information assets

HIPAA-SECURITY

- **164.312(a)(1)** — Access Control
- **164.312(d)** — Person or Entity Authentication

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
T1210	Exploitation of Remote Services	Lateral-Movement
T1550	Use Alternate Authentication Material	Defense-Evasion
T1190	Exploit Public-Facing Application	Initial-Access
T1562	Impair Defenses	Defense-Evasion
T1078	Valid Accounts	Defense-Evasion
T1110	Brute Force	Credential-Access
T1021	Remote Services	Lateral-Movement
T1068	Exploitation for Privilege Escalation	Privilege-Escalation
T1588.006	Vulnerabilities	Resource-Development

Sources

Source	URL	Tier
Blog	https://www.crowdstrike.com/en-us/blog/frontier-ai-collapses-exploi...	T3

Source	URL	Tier
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
Frontier AI for Defenders: CrowdStrike and OpenAI TAC	https://www.crowdstrike.com/en-us/blog/frontier-ai-for-defenders-cr...	T3
Anthropic and OpenAI unveil Claude Mythos and GPT-5.4-Cyber	https://www.orangecyberdefense.com/global/blog/innovation/anthropic...	T3
CrowdStrike Integrates GPT-5.4-Cyber into Falcon Platform - LinkedIn	https://www.linkedin.com/posts/getaigovernence_falcon-aiagents-aise...	T3

Generated 2026-05-01 14:03 UTC by TJS Security Command Center