INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-04-24 13:42 UTC

AI-Accelerated Exploitation Compresses Defender Response Windows, Strategic Posture Shift Required

SECURITY ANALYSIS | HIGH | CVSS 7.5

SCC Item ID	SCC-STY-2026-0081
Type	Security Analysis
Severity	HIGH
CVSS Base Score	7.5
Affected Products	Security programs broadly; CrowdStrike Falcon Platform referenced as defensive context
Discovery Source	Rss:T1 Threatintel

Executive Summary

Adversaries using frontier AI models are collapsing the time between vulnerability disclosure and active exploitation, in some observed cases to near-real-time, fundamentally breaking the assumption that defenders have days or weeks to patch before attacks begin. CrowdStrike's 2026 Global Threat Report (sources accessed via blog posts and webinar materials; direct report access recommended to verify these figures before operational use) cites 89% year-over-year growth in AI-enabled adversary activity and a 27-second observed adversary breakout time, signaling that periodic patch cycles and CVSS-backlog prioritization are no longer viable as primary defense postures. Organizations that have not shifted toward continuous monitoring, behavior-based detection, and risk-tiered response will face a structural disadvantage against adversaries who now operate at machine speed.

Technical Analysis

The central claim in CrowdStrike's 2026 Global Threat Report and associated blog series is operationally significant regardless of which specific AI models adversaries are using: adversaries using frontier AI are enabling threat actors to identify, weaponize, and deploy exploits faster than defenders can complete a single patch cycle. The reported 42% increase in zero-days exploited before public disclosure is particularly consequential - it means CVE publication and CVSS scoring, the traditional triggers for prioritization workflows, arrive after exploitation has already begun.

The 27-second adversary breakout time, the interval between initial access and lateral movement, further compresses the defensive response window. At that tempo, human-analyst-in-the-loop detection and response pipelines cannot match adversary speed. This is not a hypothetical future-state concern; it reflects observed

operational behavior per CrowdStrike's telemetry.

The MITRE ATT&CK techniques surfaced in this story trace a coherent kill chain: T1588.006 (obtaining AI capabilities as a resource), T1588 (capability acquisition broadly), T1078 (valid account abuse), T1110 (brute force for initial access), T1190 (exploit public-facing application), T1068 and T1203 (privilege escalation via exploitation), T1059 (command and scripting interpreter execution), T1210 (lateral movement via exploit), T1550 (use of alternate authentication material), and T1486 (data encrypted for impact, ransomware stage). The pattern describes adversaries using AI to accelerate capability development, then executing familiar but faster-moving intrusion chains.

Important sourcing caveat: The statistics cited (89% growth, 42% zero-day increase, 27-second breakout) are attributed to CrowdStrike's 2026 Global Threat Report. The CrowdStrike blog posts are the available sources at this time; direct report access is recommended before using these figures in operational briefings or public communications. References to 'Anthropic Claude Mythos' and 'OpenAI GPT-5.4-Cyber' as adversary tools have not been independently verified. [Note: 'Mythos' appears to be a CrowdStrike scenario-framing construct rather than a shipping Anthropic product, and 'GPT-5.4-Cyber' attributions originate from LinkedIn and Mashable rather than authoritative vendor announcements. These should not be treated as confirmed product capabilities without direct verification from OpenAI or Anthropic.] The strategic posture argument stands independent of which specific AI models are involved; the tool attribution should not be treated as confirmed without further verification.

Security programs built around quarterly scanning cadences, patch SLAs measured in weeks, and CVSS-threshold-based prioritization queues are structurally misaligned with this environment. The defensive implication is a required shift toward continuous exposure management, automated detection-and-response pipelines, and threat-intelligence-driven prioritization that does not depend on CVE publication as its trigger.

Action Checklist

1. Step 1: Assess exposure, audit whether your vulnerability management program uses CVE publication or CVSS scoring as the primary prioritization trigger; if so, that workflow is now a structural gap against pre-disclosure exploitation
2. Step 2: Review controls, verify EDR coverage completeness and behavioral detection rule quality across endpoints; confirm identity controls include MFA on all privileged accounts (T1078, T1110 are primary initial access vectors in this kill chain); review lateral movement detection capabilities given the 27-second breakout time reference
3. Step 3: Update threat model, add AI-accelerated exploitation as an active threat scenario in your risk register; map T1588.006 (AI capability acquisition) into your threat actor profiling to account for adversaries using AI for exploit development, not just delivery
4. Step 4: Communicate findings, brief leadership on the structural gap between current patch cadence and observed exploitation timelines; frame this as a program design question, not a single vulnerability to remediate
5. Step 5: Obtain CrowdStrike's 2026 Global Threat Report (if publicly available) or request a briefing directly from CrowdStrike to verify cited statistics before using them in operational or board-level communications. As of publication, the full report may be available via CrowdStrike's customer portal or public download; confirm current availability directly with CrowdStrike.

IR / Forensic Enrichment

Triage Priority	URGENT
Escalation Criteria	Escalate immediately to CISO and legal counsel if any endpoint telemetry, EDR alert, or authentication log indicates active exploitation of a vulnerability disclosed within the prior 72 hours — particularly Event ID 4648/4672 anomalies on privileged accounts or lateral movement indicators within minutes of initial access — as this pattern is consistent with AI-accelerated breakout and the 27-second timeline leaves no room for standard triage queuing; additionally escalate if MFA gaps on privileged accounts are confirmed, as this directly enables the T1078/T1110 initial access vectors cited in this threat scenario.
Recovery Notes	Because this threat scenario centers on program design failure rather than a specific active incident, recovery focuses on structural remediation: validate that your updated vulnerability management process is operationally active by running a full asset scan and confirming that CISA KEV entries published in the prior 7 days are being triaged within 24 hours rather than on the standard patch cycle. Monitor privileged account authentication logs (Windows Security Event IDs 4624, 4648, 4672) and EDR behavioral alerts for a minimum of 30 days post-remediation to confirm that MFA enforcement changes have closed the T1078/T1110 initial access paths. If the CrowdStrike GTR statistics are confirmed accurate upon direct review, update your incident response plan (per NIST IR-8) to formally encode AI-accelerated exploitation as a named threat scenario with explicit detection thresholds and escalation timelines, and re-test the updated plan within 90 days per NIST IR-3 (Incident Response Testing).
Forensic Artifacts	Windows Security Event Log — Event IDs 4624 (Successful Logon), 4625 (Failed Logon), 4648 (Explicit Credential Logon), and 4672 (Special Privilege Assigned): these are the primary artifacts of T1078 (Valid Accounts) and T1110 (Brute Force) initial access, and in an AI-accelerated scenario the 4625 failure spike will be compressed into a very short window followed almost immediately by a successful 4624, collapsing the detection window to seconds | EDR behavioral telemetry — specifically process lineage trees showing lateral movement tool execution (e.g., PsExec, WMI, SMB) within 30 seconds of an initial authentication event: the 27-second breakout time means lateral movement artifacts will appear in the same or adjacent log timestamp bucket as the initial access event, which is diagnostically distinctive from human-speed intrusions | VPN and remote access gateway authentication logs — filtered for accounts lacking MFA enforcement flags and showing successful authentication followed by no MFA challenge event: this artifact pattern directly evidences the T1078/T1110 kill chain gap and establishes which accounts were exploitable at the time of any incident | Vulnerability management scanner remediation-age report — exported as a timestamped artifact showing all open Critical and High findings with days-since-discovery: this is forensic evidence of the structural gap between patch cadence and AI-compressed exploitation timelines, and establishes organizational knowledge of unpatched exposure prior to any breach | CISA KEV catalog delta log — a dated record of all KEV entries added in the prior 30 days cross-referenced against your asset inventory: any KEV entry matching an asset in your environment that was not remediated within 24 hours is direct forensic evidence of the structural VM program gap this threat scenario describes, and would be relevant to any regulatory or cyber insurance inquiry

Per-Action IR Details

Step 1: Assess exposure — audit whether your vulnerability management program uses CVE publication or CVSS scoring as the primary prioritization trigger; if so, that workflow is now a structural gap against

pre-disclosure exploitation

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Establishing IR capability and vulnerability posture before adverse events occur

Controls: NIST SI-2 (Flaw Remediation) — specifically the requirement to identify and correct flaws before exploitation, not reactively after CVE publication, NIST RA-3 (Risk Assessment) — risk decisions must account for threat intelligence indicating pre-disclosure exploitation, not only published CVSS scores, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — process must be reviewed against AI-compressed exploitation timelines, not assumed to be adequate based on legacy patch cadence, CIS 7.2 (Establish and Maintain a Remediation Process) — remediation strategy must incorporate threat-intelligence-driven prioritization beyond CVSS thresholds

Compensating: For teams without a commercial VM platform: run weekly OpenVAS or Greenbone Community Edition scans and cross-reference output against CISA KEV (Known Exploited Vulnerabilities catalog at cisa.gov/known-exploited-vulnerabilities-catalog) as a prioritization layer that reflects active exploitation rather than CVSS alone. Supplement with a cron job or scheduled PowerShell task that pulls the CISA KEV JSON feed daily and alerts on any new entry matching assets in your inventory: `Invoke-WebRequest -Uri 'https://www.cisa.gov/sites/default/files/feeds/known_exploited_vulnerabilities.json' -OutFile kev.json` followed by asset-matching logic against your CMDB export.

Evidence: Before restructuring the VM workflow, document the current state as a baseline: export your vulnerability scanner's open findings sorted by CVE publish date versus CVSS score to quantify how many high-CVSS findings are older than 30 days with no remediation action — this establishes the structural lag that AI-accelerated exploitation exploits. Also capture your patch SLA policy document version and date, as this will be needed for any post-incident regulatory review under NIST IR-6 (Incident Reporting) obligations.

Step 2: Review controls — verify EDR coverage completeness and behavioral detection rule quality across endpoints; confirm identity controls include MFA on all privileged accounts (T1078, T1110 are primary initial access vectors in this kill chain); review lateral movement detection capabilities given the 27-second breakout time reference

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Ensuring detection tools, identity controls, and lateral movement visibility are in place before an AI-accelerated intrusion compresses dwell time to seconds

Controls: NIST SI-4 (System Monitoring) — monitoring must include behavioral detection for credential abuse (T1078 — Valid Accounts) and brute force patterns (T1110 — Brute Force), not solely signature-based alerting, NIST IA-5 (Authenticator Management) — MFA enrollment must be verified as enforced, not merely configured, on all privileged accounts given T1078 and T1110 are cited primary initial access vectors, NIST AC-2 (Account Management) — privileged account inventory must be current and auditable; a 27-second breakout time means unmanaged privileged accounts are compromised before manual review is possible, CIS 6.3 (Require MFA for Externally-Exposed Applications) — enforce MFA on all externally-exposed interfaces as a hard prerequisite given T1110 (Brute Force) is listed as a primary vector, CIS 6.5 (Require MFA for Administrative Access) — MFA on administrative accounts is non-negotiable when adversary breakout time is 27 seconds post-initial-access

Compensating: For teams without enterprise EDR: deploy Sysmon with the SwiftOnSecurity or Olaf Hartong modular config to capture process creation, network connections, and credential access events. For T1078 detection without a SIEM, run this PowerShell query on domain controllers hourly via scheduled task to flag anomalous privileged logons: `Get-WinEvent -LogName Security | Where-Object {$_.Id -eq 4624 -and $_.Message -match 'Logon Type:\s+3'} | Select-Object TimeCreated, Message | Export-Csv privileged_logons.csv`. For T1110, alert on Windows Security Event ID 4625 (Failed Logon) spikes exceeding 10 failures in 60 seconds from a single source using a simple PowerShell threshold script. Use OSQuery with the `logged_in_users` and `last` tables to audit active privileged sessions.

Evidence: Before this review, pull and preserve: (1) Windows Security Event Log entries for Event ID 4648 (Logon using explicit credentials) and 4672 (Special privileges assigned to new logon) across all domain controllers for the prior 30 days — these are the artifacts T1078 leaves at the point of privileged account abuse; (2) authentication logs

from your VPN or remote access gateway filtered for accounts with no MFA enforcement flag; (3) a CrowdStrike Falcon or equivalent EDR coverage gap report showing endpoints with sensor offline or policy exclusions — gaps here represent the blind spots AI-accelerated lateral movement will exploit given a 27-second breakout window.

Step 3: Update threat model — add AI-accelerated exploitation as an active threat scenario in your risk register; map T1588.006 (AI capability acquisition) into your threat actor profiling to account for adversaries using AI for exploit development, not just delivery

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Threat modeling and risk register updates are pre-incident activities that shape detection logic and response priorities before an incident occurs

Controls: NIST RA-3 (Risk Assessment) — the risk register must reflect current threat intelligence; treating AI-accelerated exploitation as a future risk rather than an active threat scenario is a risk assessment gap given the CrowdStrike 2026 GTR data point of 89% YoY growth in AI-enabled adversary activity, NIST RA-5 (Vulnerability Monitoring and Scanning) — vulnerability monitoring must now include threat intelligence feeds that flag pre-CVE exploitation indicators, not only post-publication scanning, NIST IR-4 (Incident Handling) — the incident handling capability must be extended to cover AI-assisted exploit development (T1588.006) as a distinct threat category with its own detection and response logic, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — the VM process must incorporate threat actor capability assessments, including AI tool acquisition (T1588.006), as a prioritization input

Compensating: For teams without a commercial threat intelligence platform: subscribe to CISA's free Automated Indicator Sharing (AIS) program and the MITRE ATT&CK STIX/TAXII feed to pull structured threat actor profiles including T1588.006 mappings. Maintain a simple threat actor matrix in a shared spreadsheet documenting which tracked groups have demonstrated AI capability acquisition, cross-referenced against your industry vertical. Use the MITRE ATT&CK Navigator (available free at mitre-attack.github.io/attack-navigator) to build a heat map of your detection coverage against the T1588.006 → T1078/T1110 → lateral movement kill chain described in this threat scenario.

Evidence: Document the current state of your risk register as a dated artifact before making updates — this establishes a before/after record demonstrating that the AI-accelerated exploitation threat scenario was formally recognized and acted upon, which is relevant to any future regulatory inquiry. Additionally, preserve any prior threat intelligence reports that characterized your adversary landscape, as these will form the baseline against which the updated threat model's improvements can be measured during post-incident review (NIST 800-61r3 §4 — Post-Incident Activity).

Step 4: Communicate findings — brief leadership on the structural gap between current patch cadence and observed exploitation timelines; frame this as a program design question, not a single vulnerability to remediate

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: Lessons learned and program-level improvements driven by threat intelligence findings feed upward to leadership to justify structural changes, even absent a specific incident

Controls: NIST IR-8 (Incident Response Plan) — the IR plan must be updated to reflect the structural gap identified; leadership briefing is the authorization step that enables plan revision and resource allocation, NIST IR-6 (Incident Reporting) — proactive reporting of structural capability gaps to leadership mirrors the incident reporting obligation and ensures decision-makers understand risk posture before an AI-accelerated breach occurs, NIST PM-9 (Risk Management Strategy) — communicating that patch cadence is structurally misaligned with AI-compressed exploitation windows is a risk management strategy communication, not a tactical finding, CIS 7.2 (Establish and Maintain a Remediation Process) — leadership briefing must result in an updated remediation process that moves from calendar-based patching to continuous, threat-intelligence-triggered remediation

Compensating: For teams without a formal GRC platform or executive reporting tool: build a one-page gap analysis document contrasting your current mean time to patch (calculate from your VM scanner's remediation age report) against the CrowdStrike-cited exploitation timeline data points (27-second breakout, near-real-time disclosure-to-exploitation in observed cases). Present this as a timeline comparison table rather than a CVSS risk

matrix — the visual contrast between 'patch in 30 days' and 'exploitation in hours' is the communication vehicle. Use CISA's free Cybersecurity Performance Goals (CPGs) as a reference framework for the recommended program design changes, as these carry regulatory credibility with leadership.

Evidence: Before the leadership briefing, collect and preserve: (1) your current documented patch SLA (e.g., Critical = 30 days, High = 90 days) as evidence of the baseline program design; (2) your VM scanner's age-of-open-findings report showing the distribution of unpatched high/critical findings by days-open — this quantifies the actual gap, not a hypothetical one; (3) any prior board or leadership communications on vulnerability management to establish continuity of the risk narrative. These artifacts protect the security team if a subsequent AI-accelerated breach occurs against a known-unpatched asset.

Step 5: Monitor developments — obtain and review CrowdStrike's 2026 Global Threat Report directly to verify cited statistics before using them in operational or board-level communications; track CISA advisories for any formal acknowledgment of AI-accelerated exploitation trends

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: Integrating current cyber threat intelligence (CTI) into ongoing monitoring and analysis is a continuous detection activity, not a one-time task

Controls: NIST SI-5 (Security Alerts, Advisories, and Directives) — CISA advisories and vendor threat reports (CrowdStrike GTR) are authoritative external sources that must be ingested, reviewed, and acted upon per this control, NIST IR-5 (Incident Monitoring) — tracking the evolution of AI-accelerated exploitation as an active threat trend is an incident monitoring activity that informs detection rule updates and response readiness, NIST AU-6 (Audit Record Review, Analysis, and Reporting) — ongoing review of external intelligence sources for AI-exploitation indicators should be formalized as a recurring analysis activity, not ad hoc consumption, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — the VM process must incorporate intelligence from CISA advisories and authoritative vendor reports as a formal input to prioritization decisions

Compensating: For teams without a commercial threat intelligence subscription: establish a free monitoring stack using (1) CISA's RSS feed for advisories (cisa.gov/news-events/cybersecurity-advisories) ingested into a free RSS reader or parsed via curl in a daily cron job; (2) a Google Alert or similar free alert configured for 'AI-accelerated exploitation CISA' and 'AI exploit development advisory'; (3) direct bookmark and quarterly review of the CrowdStrike Threat Intelligence portal for publicly released reports. For verifying cited statistics before operational use, cross-reference the GTR data points against CISA's published threat landscape documents and Mandiant's M-Trends report (also publicly available) — if two independent authoritative sources corroborate the trend, the data point is operationally usable.

Evidence: Before acting on the cited CrowdStrike statistics operationally or in board communications, document your source verification chain: record the report title, publication date, page/section of each cited statistic (89% YoY AI-enabled activity growth, 27-second breakout time), and the date you accessed and verified the source. This verification record is forensic evidence of due diligence if the statistics are later disputed or revised, and it satisfies the NIST AU-10 (Non-Repudiation) requirement for accountability in decision-making. Also preserve any CISA advisory URLs and access dates for the same reason.

Detection Guidance

Given the compressed breakout time and pre-disclosure exploitation pattern, detection focus should shift from signature-based CVE matching toward behavioral anomaly detection. Key areas to monitor:

Identity and authentication: Log and alert on unusual authentication patterns consistent with T1078 (valid account abuse) and T1550 (pass-the-hash, pass-the-ticket), specifically, lateral authentication from accounts that do not normally move across systems, and authentication outside normal business hours or from atypical source IPs.

Initial access via exploitation: Monitor web-facing application logs for exploitation attempt patterns (T1190) independent of known CVE signatures; behavior-based WAF and application telemetry are more reliable than

signature matching when exploits precede disclosure.

Privilege escalation: Alert on process execution chains consistent with T1068 and T1203, specifically, user-space processes spawning elevated child processes without an expected workflow trigger.

Lateral movement velocity: Given the 27-second breakout reference, configure SIEM rules to flag rapid lateral authentication attempts within short time windows across multiple hosts. SMB and WMI lateral movement patterns (T1210) should be baselined and monitored continuously, not periodically.

Scripting and execution: T1059 (command and scripting interpreter) detections should be reviewed for execution from unusual parent processes or unexpected user contexts.

Hunting hypothesis: Search for accounts that authenticated successfully, then within 60 seconds accessed additional internal systems, this pattern is consistent with automated lateral movement at machine speed.

Note: No specific file hashes, C2 domains, or IP indicators are available from the provided source material for this story. IOC section reflects this accurately.

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 referenced as the source for AI-enabled adversary activity telemetry; specific IOCs including tool signatures, C2 infrastructure, or payload hashes associated with AI-accelerated exploitation campaigns are not present in the available source blog posts and require direct report access	LOW

Framework Mappings

MITRE-ATTACK

- **T1588.006** — Vulnerabilities
- **T1588** — Obtain Capabilities
- **T1078** — Valid Accounts
- **T1110** — Brute Force
- **T1486** — Data Encrypted for Impact
- **T1059** — Command and Scripting Interpreter
- **T1190** — Exploit Public-Facing Application
- **T1068** — Exploitation for Privilege Escalation
- **T1203** — Exploitation for Client Execution
- **T1210** — Exploitation of Remote Services
- **T1550** — Use Alternate Authentication Material

NIST-800-53R5

- **AC-2** — Account Management
- **AC-6** — Least Privilege
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **AC-7** — Unsuccessful Logon Attempts
- **CP-9** — System Backup
- **CP-10** — System Recovery and Reconstitution
- **CM-7** — Least Functionality
- **SI-3** — Malicious Code Protection
- **SI-4** — System Monitoring
- **SI-7** — Software, Firmware, and Information Integrity
- **CA-8** — Penetration Testing
- **RA-5** — Vulnerability Monitoring and Scanning
- **SC-7** — Boundary Protection
- **SI-2** — Flaw Remediation
- **AC-3** — Access Enforcement
- **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

- **3.3** — Configure Data Access Control Lists
- **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
- **6.4** — Require MFA for Remote Network Access
- **6.5** — Require MFA for Administrative Access
- **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(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
T1588.006	Vulnerabilities	Resource-Development
T1588	Obtain Capabilities	Resource-Development
T1078	Valid Accounts	Defense-Evasion
T1110	Brute Force	Credential-Access
T1486	Data Encrypted for Impact	Impact
T1059	Command and Scripting Interpreter	Execution
T1190	Exploit Public-Facing Application	Initial-Access
T1068	Exploitation for Privilege Escalation	Privilege-Escalation
T1203	Exploitation for Client Execution	Execution
T1210	Exploitation of Remote Services	Lateral-Movement
T1550	Use Alternate Authentication Material	Defense-Evasion

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
CrowdStrike Integrates GPT-5.4-Cyber into Falcon Platform - LinkedIn	https://www.linkedin.com/posts/getaigovernence_falcon-aiagents-aise...	T3

Source	URL	Tier
OpenAI follows Anthropic's lead in limited release of GPT■5.4■Cyber	https://mashable.com/article/openai-gpt-54-cyber-cybersecurity-ai-m...	T3

Generated 2026-04-24 13:42 UTC by TJS Security Command Center