INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-05-17 13:50 UTC

# AI-Accelerated Vulnerability Discovery Is Rewriting the Patching Playbook, and Defenders Are Behind

SECURITY ANALYSIS | HIGH | CVSS 7.5

| | |
|---|---|
| SCC Item ID | SCC-STY-2026-0139 |
| Type | Security Analysis |
| Severity | HIGH |
| CVSS Base Score | 7.5 |
| Affected Products | Enterprise environments broadly; CrowdStrike Falcon Platform (integrating Claude Opus 4.7); organizations relying on CVSS-driven patch cycles and static asset inventories |
| Discovery Source | Rss:T1 Threatintel |

## Executive Summary

Frontier AI models are now discovering zero-day vulnerabilities faster than enterprise patch cycles can respond, collapsing the window between discovery and exploitation that defenders have historically relied on. CrowdStrike's 2026 Global Threat Report documents a 42% year-over-year increase in zero-days exploited before public disclosure, and adversaries are actively fine-tuning offensive AI models on tailored attack datasets, a shift from human-speed to machine-speed threat operations. For security leaders, this signals that monthly patch cadences and static asset inventories are no longer viable as primary risk controls; continuous exposure management and real-time detection coverage must replace them.

## Technical Analysis

The structural risk introduced by AI-accelerated vulnerability discovery is not theoretical. Three converging trends define the current threat posture.

First, the discovery-to-exploitation timeline has compressed. AI-assisted fuzzing, automated code analysis, and vulnerability chaining, all well-documented in MITRE ATT&CK and academic offensive research, now operate at speeds that outpace the average enterprise patch window. When adversaries weaponize a zero-day before public disclosure, the standard CVSS-driven triage model fails: defenders are triaging a vulnerability that has already been exploited.

Second, adversaries are industrializing offensive AI capability. IBM X-Force reporting indicates threat actors are constructing tailored datasets to fine-tune attack-focused models, effectively building bespoke reconnaissance and exploitation tooling. This operationalization of offensive AI maps broadly to MITRE ATT&CK T1588 (Obtain

Capabilities), T1595 (Active Scanning), T1587.001 (Develop Capabilities: Malware), and T1583 (Acquire Infrastructure). Threat actors including APT28 (FANCY BEAR), Lazarus Group (APT38), and other state-linked actors have demonstrated the operational sophistication to adopt and adapt emerging offensive tooling at scale. Consult MITRE ATT&CK threat actor profiles for current TTP mappings.

Third, the vulnerability classes most susceptible to automated discovery, memory safety errors (CWE-119: Buffer Errors, CWE-787: Out-of-Bounds Write, CWE-416: Use After Free) and input validation failures (CWE-20), represent a substantial share of the existing unpatched attack surface in enterprise environments. These are not novel weakness types; they are familiar weaknesses newly discoverable at machine speed. (This is a strategic intelligence item, not a specific CVE; CVSS scoring is not applicable.)

On the defender side, CrowdStrike announced integration of an advanced Anthropic Claude model variant (referenced as Claude Opus 4.7 in press materials) into the Falcon platform. The specific model version and formal structure of Project Glasswing - described as including Anthropic, CrowdStrike, Microsoft, Google, and Apple - could not be independently verified against public vendor disclosures at the time of this analysis; that framing should be treated as low-confidence until confirmed through primary source publication.

The operational implication for security teams is direct: CVSS score alone cannot drive prioritization when exploitation may precede NVD or CISA KEV publication. Detection engineering must account for T1190 (Exploit Public-Facing Application) and T1068 (Exploitation for Privilege Escalation) occurring against vulnerabilities with no published CVE at triage time. Hunting hypotheses should target anomalous code execution patterns, unexpected process chains from internet-facing services, and privilege escalation activity that precedes any patch advisory.

## Action Checklist

1. Assess exposure, audit all internet-facing and externally accessible assets for known instances of CWE-119, CWE-787, CWE-416, and CWE-20 weakness classes in deployed software; prioritize applications written in memory-unsafe languages without compensating controls
2. Review patch cadence, evaluate whether your current patch cycle (monthly or longer) is defensible given AI-accelerated exploitation timelines; identify which asset classes require continuous patching or virtual patching via WAF or EDR policy rather than scheduled maintenance windows
3. Validate detection coverage for pre-disclosure exploitation, confirm EDR and SIEM rules cover T1190, T1068, T1203, and T1210 for assets where a zero-day could plausibly exist; ensure detection logic does not depend on a CVE being published before alerting
4. Stress-test asset inventory, verify that your asset inventory is dynamic and continuously reconciled, not static; AI-assisted scanning (T1595) by adversaries will discover assets your inventory does not know exist
5. Update threat model for sophisticated state-linked actors (e.g., APT28/FANCY BEAR, APT38/Lazarus Group), map documented TTPs against your current control environment; where gaps exist against reconnaissance, capability development, or infrastructure staging, document and prioritize remediation. Consult MITRE ATT&CK threat actor profiles for current TTP mappings.
6. Brief leadership, communicate that the 42% increase in pre-disclosure zero-day exploitation (CrowdStrike 2026 Global Threat Report, medium confidence) represents a structural change in attacker capability, not a temporary spike; frame the ask around continuous exposure management investment rather than incremental patching improvements

7. Monitor Anthropic and participating vendor public disclosures regarding Project Glasswing and AI-accelerated vulnerability detection capabilities; if formally announced, reassess your AI-assisted defense roadmap. Do not delay execution of items 1-6 while awaiting confirmation.

## IR / Forensic Enrichment

| | |
|---|---|
| **Triage Priority** | URGENT |
| **Escalation Criteria** | Escalate immediately to CISO and legal counsel if forensic evidence from Steps 3 or 4 reveals active pre-disclosure exploitation attempts against memory-unsafe internet-facing services (behavioral T1190/T1068 indicators firing without a corresponding published CVE), or if asset reconciliation (Step 4) discovers internet-exposed assets processing PII/PHI/PCI data that are absent from the current inventory, triggering potential breach notification obligations under HIPAA, GDPR, or applicable state law. |
| **Recovery Notes** | Post-containment, prioritize re-validation of all internet-facing services for CWE-119/787/416/20 class weaknesses using updated SCA and binary analysis tooling, as AI-accelerated adversary discovery means previously unknown weaknesses in your stack may have been catalogued by threat actors before your own assessment completed. Monitor behavioral detection rules (T1190, T1068, T1203, T1210) with elevated sensitivity for a minimum of 30 days post-remediation, as FANCY BEAR and FAMOUS CHOLLIMA are known to maintain persistent access and re-enter environments after defenders believe threats are eradicated. Conduct a lessons-learned session within 5 business days focused specifically on whether your patch cadence policy change (Step 2) is operationally implemented, not just documented — the structural attacker capability shift described in the CrowdStrike 2026 Global Threat Report makes policy-without-execution the primary residual risk. |

| | |
|---|---|
| **Forensic Artifacts** | Web server and reverse proxy access logs (Apache access.log, Nginx access.log, IIS W3C logs) covering 30+ days prior to assessment — AI-assisted T1595 reconnaissance and T1190 exploitation of CWE-119/787/416/20-class vulnerabilities produce distinctive anomalous HTTP request patterns (oversized headers, binary payloads in GET parameters, unexpected URI paths) that predate CVE publication and are recoverable from these logs if not rotated \| Sysmon Event ID 1 (Process Creation) and Event ID 10 (Process Access) logs from internet-facing Windows hosts — T1068 exploitation of memory-unsafe services will produce child process chains where the vulnerable service (e.g., IIS worker w3wp.exe or a C/C++ daemon) spawns cmd.exe, powershell.exe, or a network utility without a corresponding user session, which Sysmon captures with parent-child PID relationships and command-line arguments \| Linux auditd execve and open syscall records for internet-facing hosts — buffer overflow exploitation (CWE-119/787) and use-after-free (CWE-416) attacks against memory-unsafe daemons generate abnormal execve chains from the vulnerable process UID/GID, recoverable from /var/log/audit/audit.log if auditd is configured with '-a always,exit -F arch=b64 -S execve' rules scoped to the service account \| Firewall NetFlow or connection state logs covering the prior 30–90 days — AI-assisted T1595 scanning by FANCY BEAR, FAMOUS CHOLLIMA, and PUNK SPIDER produces high-volume, sequential port-sweep patterns from rotating source IPs with sub-second inter-packet timing that is statistically distinguishable from human-driven reconnaissance in flow records and can be recovered from perimeter firewall or cloud security group flow logs \| EDR process tree telemetry and memory snapshot artifacts for any process flagged by behavioral T1068/T1190 rules — if CrowdStrike Falcon (referenced in the threat context as integrating Claude Opus 4.7) is deployed, pull the full process tree dump and any associated memory scan artifacts from the Falcon console for the vulnerable service processes; for non-EDR environments, use Volatility3 (free, open-source) against a live memory image captured with WinPmem or LiME to recover injected shellcode or heap spray artifacts characteristic of CWE-787 (out-of-bounds write) exploitation |

### Per-Action IR Details

**Assess exposure — audit all internet-facing and externally accessible assets for known instances of CWE-119, CWE-787, CWE-416, and CWE-20 weakness classes in deployed software; prioritize applications written in memory-unsafe languages without compensating controls**

**NIST Phase:** Preparation

**Reference:** NIST 800-61r3 §2 — Preparation: Establishing IR Capability and Asset Visibility

**Controls:** NIST SI-2 (Flaw Remediation), NIST RA-3 (Risk Assessment), NIST CM-8 (System Component Inventory), CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory), CIS 7.1 (Establish and Maintain a Vulnerability Management Process)

**Compensating:** Run Grype or Trivy (both free, open-source SCA tools) against container images and package manifests to flag CWE-119/787/416/20-class libraries — e.g., 'grype dir:. --output json | jq .matches[].vulnerability.id'. For host-level C/C++ binaries, use checksec (free) to identify binaries lacking stack canaries, NX, or PIE: 'checksec --format=json --file=/usr/bin/target_binary'. Cross-reference output against your asset list manually in a spreadsheet; a 2-person team can triage by exposure tier (internet-facing first).

**Evidence:** Before remediating, snapshot the current software bill of materials (SBOM) for all internet-facing services as a baseline — capture 'dpkg -l > sbom_baseline_$(hostname)_$(date +%F).txt' or equivalent. Preserve existing WAF and load balancer access logs covering the prior 30 days, as AI-assisted reconnaissance (T1595.002 — Active Scanning: Vulnerability Scanning) against these weakness classes may already be visible as anomalous payload patterns in HTTP request bodies targeting memory-unsafe endpoints.

**Review patch cadence — evaluate whether your current patch cycle (monthly or longer) is defensible given AI-accelerated exploitation timelines; identify which asset classes require continuous patching or virtual**

**patching via WAF or EDR policy rather than scheduled maintenance windows**

**NIST Phase:** Preparation

**Reference:** NIST 800-61r3 §2 — Preparation: Policy and Procedure Readiness

**Controls:** NIST SI-2 (Flaw Remediation), NIST SI-5 (Security Alerts, Advisories, and Directives), NIST CM-3 (Configuration Change Control), 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:** For teams without an enterprise patch management platform: use unattended-upgrades (Linux) with 'Unattended-Upgrade::Allowed-Origins' scoped to security repos only, or configure Windows Update via GPO to 'Auto download and install' for security classifications. For virtual patching without a commercial WAF, deploy ModSecurity (free, open-source) with the OWASP Core Rule Set (CRS) — enable paranoia level 2+ for internet-facing apps processing untrusted input, which will block common CWE-20 (improper input validation) exploitation attempts without a patch being available.

**Evidence:** Document current patch SLAs per asset class in writing before this review — this creates an auditable baseline against which the post-review policy change is measured (supports NIST IR-8 IR Plan documentation). Pull the last 90 days of patch deployment timestamps from WSUS, yum history, or apt logs to quantify actual patch lag per asset tier: 'rpm -qa --queryformat "%{installtime:date} %{name}-%{version}\n" | sort' or 'grep "upgraded" /var/log/dpkg.log | tail -200'. This evidence establishes whether current cadence already failed to meet any existing internal SLA.

**Validate detection coverage for pre-disclosure exploitation — confirm EDR and SIEM rules cover T1190, T1068, T1203, and T1210 for assets where a zero-day could plausibly exist; ensure detection logic does not depend on a CVE being published before alerting**

**NIST Phase:** Detection Analysis

**Reference:** NIST 800-61r3 §3.2 — Detection and Analysis: Identifying Adverse Events Without Prior Signature

**Controls:** NIST SI-4 (System Monitoring), NIST AU-6 (Audit Record Review, Analysis, and Reporting), NIST IR-4 (Incident Handling), CIS 8.2 (Collect Audit Logs)

**Compensating:** Deploy Sysmon (free, Microsoft Sysinternals) with the SwiftOnSecurity or Olaf Hartong modular config — this gives process creation (Event ID 1), network connection (Event ID 3), and process injection (Event ID 8/10) telemetry that detects T1068 (privilege escalation via exploit) and T1190 (initial access via public-facing app) behaviorally, independent of CVE publication. Load the Sigma rule set (free, community-maintained) for these techniques into your log pipeline using sigma-cli: 'sigma convert -t splunk rules/windows/process_creation/proc_creation_win_exploit_*.yml'. For T1203 (client-side exploit): enable PowerShell ScriptBlock logging (Event ID 4104) and Sysmon Event ID 1 filtering for Office/browser child processes spawning cmd.exe or powershell.exe.

**Evidence:** Before tuning detection rules, export and preserve the current state of all active SIEM/EDR detection rules as a snapshot — this is your pre-improvement baseline for post-incident audit. Query existing logs NOW for behavioral indicators of T1190 exploitation attempts against memory-unsafe services: look for web server worker processes (w3wp.exe, nginx, apache2) spawning unexpected child processes in Windows Security Event Log Event ID 4688 (Process Creation) or Linux auditd execve syscalls. For T1068, search for privilege escalation sequences: process running as low-privilege user followed within seconds by a child process running as SYSTEM/root without an expected sudo/UAC event.

**Stress-test asset inventory — verify that your asset inventory is dynamic and continuously reconciled, not static; AI-assisted scanning (T1595) by adversaries will discover assets your inventory does not know exist**

**NIST Phase:** Preparation

**Reference:** NIST 800-61r3 §2 — Preparation: Understanding the Environment Before an Incident Occurs

**Controls:** NIST CM-8 (System Component Inventory), NIST CA-7 (Continuous Monitoring), NIST RA-3 (Risk Assessment), CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory), CIS 1.2 (Address Unauthorized Assets)

**Compensating:** Run nmap with OS and service detection against your own declared IP ranges on a weekly cron job and diff the output against your last known-good inventory: 'nmap -sV -O --open -oX scan_$(date +%F).xml 10.0.0.0/8 && ndiff scan_last.xml scan_$(date +%F).xml > delta_$(date +%F).txt'. Use osquery (free) with a scheduled query on 'listening_ports' and 'interface_addresses' tables to detect new services on existing hosts that weren't present in the prior week's baseline. Flag any delta as requiring immediate classification — adversaries using AI-assisted T1595 reconnaissance will find these assets within hours of exposure.

**Evidence:** Before running your own reconciliation scan, pull DNS query logs from your recursive resolvers covering the past 30 days and look for lookups of internal hostnames from external sources or lookups of shadow IT domains — this may reveal assets already being probed by T1595.001 (Active Scanning: Scanning IP Blocks). Preserve firewall flow logs (NetFlow/IPFIX or iptables log entries) for the same period; AI-assisted scanners generate characteristically high-volume, low-dwell connection attempts across broad port ranges that are distinguishable from normal traffic patterns in flow data.

**Update threat model for FANCY BEAR, FAMOUS CHOLLIMA, and PUNK SPIDER — map documented TTPs for these actors against your current control environment; where gaps exist against T1588.006, T1587.001, or T1583, document and prioritize remediation**

**NIST Phase:** Preparation

**Reference:** NIST 800-61r3 §2 — Preparation: Threat Intelligence Integration and Scenario Planning

**Controls:** NIST RA-3 (Risk Assessment), NIST PM-16 (Threat Awareness Program), NIST IR-8 (Incident Response Plan), CIS 7.1 (Establish and Maintain a Vulnerability Management Process)

**Compensating:** Use the free MITRE ATT&CK Navigator (web-based, no install required at https://mitre-attack.github.io/attack-navigator/) to layer FANCY BEAR (G0007), FAMOUS CHOLLIMA (G1027), and PUNK SPIDER (G0092) technique heatmaps against your control coverage — export as JSON and use as a living gap register. For T1588.006 (Acquire Access: Vulnerability Databases) and T1587.001 (Develop Capabilities: Malware), monitor GitHub, Exploit-DB, and Packet Storm RSS feeds using a free RSS aggregator or a simple Python script with feedparser to alert on PoC code referencing your deployed software stack. This gives a 2-person team early warning of adversary tooling development without a CTI subscription.

**Evidence:** Before updating the threat model, preserve the current threat model document with a date stamp as evidence of the pre-improvement state — critical for demonstrating due diligence in regulatory or legal contexts. Pull 90 days of outbound DNS and proxy logs and search for domains associated with FANCY BEAR C2 infrastructure (MITRE G0007 references Sofacy/X-Agent C2 patterns) and FAMOUS CHOLLIMA tooling (associated with DPRK-nexus spearphishing and supply chain staging domains); use abuse.ch URLhaus (free) blocklist exports as a reference feed for known-bad domains attributable to these actors.

**Brief leadership — communicate that the 42% increase in pre-disclosure zero-day exploitation (CrowdStrike 2026 Global Threat Report, medium confidence) represents a structural change in attacker capability, not a temporary spike; frame the ask around continuous exposure management investment rather than incremental patching improvements**

**NIST Phase:** Post Incident

**Reference:** NIST 800-61r3 §4 — Post-Incident Activity: Lessons Learned and Strategic Communication

**Controls:** NIST IR-6 (Incident Reporting), NIST IR-8 (Incident Response Plan), NIST PM-9 (Risk Management Strategy), CIS 7.1 (Establish and Maintain a Vulnerability Management Process), CIS 7.2 (Establish and Maintain a Remediation Process)

**Compensating:** A 2-person team can build a credible executive brief using free resources: pull your own patch SLA data (from Step 2 evidence) and overlay it against the CrowdStrike-reported 42% pre-disclosure exploitation increase to show the gap quantitatively. Use CISA's Known Exploited Vulnerabilities (KEV) catalog (free, CSV download at cisa.gov/known-exploited-vulnerabilities-catalog) to calculate how many KEV entries affecting your stack had a public PoC available within 7 days of CVE publication — this is concrete evidence that your current monthly cycle is structurally misaligned with attacker timelines, and requires no paid tooling to produce.

**Evidence:** Before the leadership brief, document in writing (with timestamps) the specific control gaps identified in Steps 1–5 — this creates an auditable record that leadership was informed of material risk, which is relevant to NIST

IR-6 (Incident Reporting) obligations and potential regulatory disclosure requirements if a subsequent incident occurs. Preserve the CrowdStrike 2026 Global Threat Report citation and confidence level (medium) in the brief itself so leadership understands the evidentiary basis and does not treat it as confirmed fact; note the 42% figure should be corroborated against CISA or other primary sources before treating as definitive.

**Monitor for Project Glasswing and Anthropic public disclosures — if the coalition and AI capability claims are formally confirmed, reassess your AI-assisted defense tooling roadmap accordingly; do not act on the low-confidence coalition framing until primary sources publish**

**NIST Phase:** Detection Analysis

**Reference:** NIST 800-61r3 §3.2 — Detection and Analysis: Integrating Cyber Threat Intelligence into Adverse Event Analysis (DE.AE-07)

**Controls:** NIST SI-5 (Security Alerts, Advisories, and Directives), NIST RA-3 (Risk Assessment), NIST IR-5 (Incident Monitoring), CIS 7.1 (Establish and Maintain a Vulnerability Management Process)

**Compensating:** Set up free monitoring for authoritative disclosures using Google Alerts (no cost) scoped to 'Project Glasswing site:anthropic.com OR site:crowdstrike.com OR site:cisa.gov' and CISA's RSS feed for advisories (https://www.cisa.gov/cybersecurity-advisories — RSS available). Create a simple tracking entry in your threat intelligence register (even a dated spreadsheet row) that records: source, confidence level (currently low), trigger condition for re-evaluation, and assigned owner — this satisfies NIST IR-5 (Incident Monitoring) documentation requirements without any paid tooling.

**Evidence:** Before acting on any future confirmed disclosure, preserve the current state of your AI-assisted defense tooling inventory and roadmap documentation — if Project Glasswing is confirmed and materially changes the capability landscape (e.g., Claude Opus 4.7 integration in CrowdStrike Falcon), you will need a dated baseline to demonstrate that your roadmap was updated in response to a confirmed disclosure rather than speculation. Flag this monitoring task in your threat intelligence register with a confidence annotation of 'low — do not operationalize' until Anthropic or CrowdStrike publish primary source confirmation, consistent with the action step's own guidance.

## Detection Guidance

Detection for AI-accelerated zero-day exploitation must shift left of CVE publication. Security teams should hunt for behavioral indicators rather than waiting for signature updates tied to disclosed vulnerabilities.

Log sources to prioritize: web application and API gateway logs for unexpected input patterns or anomalous parameter structures targeting memory boundaries; EDR telemetry for process creation chains spawned from internet-facing services (T1190, T1203); privilege escalation events (T1068) on hosts where no patch advisory exists; and network flow data for scanning patterns (T1595) that suggest automated, high-speed enumeration rather than human-paced reconnaissance.

Behavioral anomalies to hunt: unusual child processes spawned from web servers, application pools, or containerized services; heap spray or memory corruption indicators in application crash telemetry; unexpected outbound connections from previously quiet internal hosts following exploitation of an internet-facing service (T1210); and new scheduled tasks or persistence mechanisms linked to infrastructure staging appearing within hours of a scanning event.

Policy gaps to audit: confirm that virtual patching rules (WAF, RASP, or EDR exploit mitigation policies) are active for all internet-facing services, not just those with published CVEs; verify that memory-safe language migration or compiler-level mitigations (stack canaries, ASLR, CFG) are enforced on highest-risk application tiers; review whether CVSS thresholds in your vulnerability management program still reflect actual exploitation probability given pre-disclosure exploitation rates.

For organizations using CrowdStrike Falcon with advanced Claude model integration, review AI-assisted triage output for coverage of CWE-119, CWE-787, CWE-416, and CWE-20 weakness classes in your asset inventory.

## 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 documents threat actor TTPs associated with AI-assisted vulnerability exploitation campaigns; specific IOC values (hashes, C2 infrastructure, tooling artifacts) are not reproduced in the source material provided for this analysis | LOW |
| TOOL | `Pending — refer to IBM X-Force reporting on offensive AI tooling` | IBM X-Force documents adversary use of fine-tuned offensive AI models for vulnerability discovery and exploitation; specific tooling indicators are not available in the source material provided for this analysis | LOW |

## Framework Mappings

### MITRE-ATTACK

- **T1588.006** — Vulnerabilities
- **T1595** — Active Scanning
- **T1587.001** — Malware
- **T1190** — Exploit Public-Facing Application
- **T1203** — Exploitation for Client Execution
- **T1068** — Exploitation for Privilege Escalation
- **T1583** — Acquire Infrastructure
- **T1210** — Exploitation of Remote Services

### NIST-800-53R5

- **CA-7** — Continuous Monitoring
- **SC-7** — Boundary Protection
- **SI-4** — System Monitoring
- **CA-8** — Penetration Testing
- **RA-5** — Vulnerability Monitoring and Scanning
- **SI-2** — Flaw Remediation
- **SI-7** — Software, Firmware, and Information Integrity
- **SI-3** — Malicious Code Protection

- **AC-6** — Least Privilege
- **SI-16** — Memory Protection
- **SI-10** — Information Input Validation
- **IR-5** — Incident Monitoring

**OWASP-TOP10-2021**

- **A03:2021** — Injection

**CIS-V8**

- **16.10** — Apply Secure Design Principles in Application Architectures
- **7.3** — Perform Automated Operating System Patch Management
- **7.4** — Perform Automated Application Patch Management

**ISO-27001-2022**

- **A.8.26** — Application security requirements
- **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 |
| T1595 | Active Scanning | Reconnaissance |
| T1587.001 | Malware | Resource-Development |
| T1190 | Exploit Public-Facing Application | Initial-Access |
| T1203 | Exploitation for Client Execution | Execution |
| T1068 | Exploitation for Privilege Escalation | Privilege-Escalation |
| T1583 | Acquire Infrastructure | Resource-Development |
| T1210 | Exploitation of Remote Services | Lateral-Movement |

## Sources

| Source | URL | Tier |
|---|---|---|
| **Blog** | https://www.crowdstrike.com/en-us/blog/tune-in-future-of-ai-powered... | T3 |
| | https://www.crowdstrike.com/en-us/blog/crowdstrike-named-leader-gar... | T3 |

| Source | URL | Tier |
|---|---|---|
| | https://www.forrester.com/blogs/project-glasswing-shows-that-ai-wil... | T2 |
| | https://www.ibm.com/think/x-force/understanding-future-of-offensive... | T3 |
| **CrowdStrike Puts Claude Opus 4.7 to Work Across Falcon and ...** | https://www.crowdstrike.com/en-us/press-releases/crowdstrike-puts-c... | T3 |



Generated 2026-05-17 13:50 UTC by TJS Security Command Center