INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-04-25 18:38 UTC

# Claude Mythos Preview Redraws the Vulnerability Discovery Threat Line: What SOC Teams Must Do Before August

SECURITY ANALYSIS | HIGH | CVSS 7.5

| | |
|---|---|
| SCC Item ID | SCC-STY-2026-0085 |
| Type | Security Analysis |
| Severity | HIGH |
| CVSS Base Score | 7.5 |
| Affected Products | Anthropic Claude Mythos Preview (frontier AI model); CrowdStrike Falcon Platform, Charlotte AI, Falcon AIDR, Falcon Data Security, AgentWorks; major OS and browser ecosystems (unspecified versions) |
| Discovery Source | Rss:T1 Threatintel |

## Executive Summary

Anthropic has reportedly demonstrated autonomous discovery of thousands of previously unknown vulnerabilities (claims unverified against primary source documentation) across major operating systems and browsers, including flaws that reportedly persisted undetected for decades despite extensive human and automated review. The development prompted the formation of Project Glasswing, a coalition of 12 major technology vendors including CrowdStrike, AWS, Apple, Google, and Microsoft, to accelerate defensive AI deployment before the model or comparable systems reach adversarial actors. This marks a structural shift in the vulnerability discovery threat model: the asymmetry between offense and defense may compress dramatically if frontier AI capability reaches state-sponsored actors or sophisticated criminal groups before enterprise defenses mature.

## Technical Analysis

Anthropic's Claude Mythos Preview represents a reported qualitative leap in automated vulnerability research. According to CrowdStrike's blog (crowdstrike.com/en-us/blog/crowdstrike-founding-member-anthropic-mythos-frontier-model-to-secure-ai/) and Anthropic's Glasswing coalition page (anthropic.com/glasswing), the model autonomously identified thousands of previously unknown vulnerabilities across major OS and browser codebases, including flaws that survived 27 years of human code review and five million automated test executions. These figures are sourced from vendor reporting and have not been independently verified against primary Anthropic publications; security teams

should treat them as directional, not precise, until primary source documentation is reviewed.

The CWE categories associated with this story reflect the classes of vulnerabilities the model reportedly discovers: CWE-119 (buffer errors, a foundational memory safety class), CWE-416 (use-after-free), CWE-269 (improper privilege management), CWE-94 (code injection), and CWE-400 (uncontrolled resource consumption). These map directly to high-impact exploit primitives. MITRE ATT&CK techniques T1068 (exploitation for privilege escalation), T1190 (exploit public-facing application), T1203 (exploitation for client execution), T1587.001 (develop capabilities: malware), T1588.006 (obtain capabilities: AI tools), T1195.001 (supply chain compromise: compromise software dependencies), and T1059 (command and scripting interpreter) describe the adversarial kill chain that becomes more accessible if this discovery capability reaches threat actors.

Project Glasswing's stated threat model explicitly names state-sponsored actors from China, Iran, North Korea, and Russia, as well as AI-augmented criminal actors identified in CrowdStrike's 2026 Global Threat Report. The coalition's formation logic is essentially a defensive race: get frontier vulnerability discovery capability embedded in detection and response tooling before adversaries operationalize it for offensive use.

CrowdStrike is integrating Mythos capabilities across Charlotte AI, Falcon AIDR, Falcon Data Security, and AgentWorks. The practical security implication for enterprise teams is twofold. First, the attack surface for AI agent systems themselves becomes a priority risk domain: if these integrations are misconfigured or exposed, they represent high-value targets. Second, the vulnerability classes Mythos discovers are the same classes that underpin the most impactful exploits in the current threat landscape. Organizations running unpatched legacy OS or browser infrastructure face elevated risk as this capability proliferates.

The EU AI Act's August 2, 2026 enforcement phase adds a compliance dimension. High-capability AI systems will face new governance requirements, and enterprise security teams deploying or integrating frontier models in security tooling need to understand both the technical exposure and the regulatory posture before that deadline.

## Action Checklist

1. Step 1: Assess AI agent exposure, inventory all deployed AI agents, LLM integrations, and agentic security tooling across your environment, including CrowdStrike Charlotte AI, Falcon AIDR, and any third-party AI-assisted security platforms
2. Step 2: Review patch posture on affected vulnerability classes, prioritize patching for CWE-119 (buffer errors), CWE-416 (use-after-free), and CWE-269 (privilege escalation) across OS and browser components, as these map to the discovery classes attributed to Mythos
3. Step 3: Update threat model, add 'AI-augmented autonomous vulnerability discovery by state-sponsored or criminal actors' as an explicit threat scenario in your threat register, mapped to T1587.001, T1588.006, and T1068
4. Step 4: Audit AI governance posture against EU AI Act timeline, if your organization deploys or integrates high-capability AI in security operations, assign ownership for August 2, 2026 compliance readiness now; engage legal and compliance counsel on classification of AI tools under the Act
5. Step 5: Monitor Project Glasswing and Anthropic primary sources directly. Verify specific quantitative claims (vulnerability counts, test execution figures) in Anthropic's technical publications and CrowdStrike's detailed blog post before using them in internal risk communications. Note that the Anthropic Glasswing coalition page (anthropic.com/glasswing) announces the initiative but may not contain detailed quantitative breakdown; check CrowdStrike's blog for specific claims.

## IR / Forensic Enrichment

| | |
|---|---|
| **Triage Priority** | URGENT |
| **Escalation Criteria** | Escalate to CISO and legal counsel immediately if: (1) any endpoint detection (Falcon AIDR, Charlotte AI alert, or Sysmon Event ID 1 process creation) shows exploitation of a CWE-416 or CWE-119 vulnerability in a browser or OS component on a host with access to sensitive data or privileged credentials; (2) any AI agent or LLM integration in your environment is confirmed to have received or processed externally-sourced prompts that enumerate internal vulnerability surfaces (potential AI-augmented reconnaissance, mapped to T1588.006); or (3) your organization is subject to EU AI Act High-Risk classification obligations and has not assigned compliance ownership before the August 2, 2026 deadline, triggering a regulatory exposure that requires legal counsel engagement. |
| **Recovery Notes** | Post-containment, verify patch integrity on all OS and browser components in the CWE-119/416/269 classes by re-running your vulnerability scanner against previously patched hosts and confirming zero residual findings in those CWE categories — do not rely solely on patch management confirmation logs, as AI-augmented attackers may target the patch verification gap. Monitor Windows Security Event Log Event ID 4688 (Process Creation) and Linux auditd execve syscall logs for at least 30 days post-patching for any anomalous child processes spawned by browser or OS subsystem processes, which would indicate either pre-patch exploitation persistence or exploitation of a newly-disclosed Glasswing vulnerability. Maintain a watching brief on Anthropic and CrowdStrike Glasswing coalition disclosures for 90 days, as coalition members have indicated they will publish integration specifics incrementally — each new disclosure may require a patch posture reassessment cycle. |
| **Forensic Artifacts** | Windows Application and System Event Logs filtered for crash/hang events (Event ID 1000, 1001, 1002) on chrome.exe, msedge.exe, firefox.exe, and ntoskrnl.exe — UAF (CWE-416) and buffer error (CWE-119) exploitation attempts frequently produce heap corruption crashes (exception code 0xC0000374) or access violations (0xC0000005) before a stable exploit is achieved \| CrowdStrike Falcon process tree telemetry for Charlotte AI and Falcon AIDR agent processes — if an AI agent has been compromised or manipulated via prompt injection as part of AI-augmented reconnaissance (T1588.006), the process tree will show anomalous child process spawning or unexpected outbound API calls from the agent process \| Linux kernel logs (/var/log/kern.log, 'dmesg' output) filtered for 'segfault', 'general protection fault', 'BUG: unable to handle kernel NULL pointer dereference' — these kernel-space fault signatures are consistent with CWE-119 buffer error exploitation attempts against kernel components in the vulnerability classes attributed to Mythos discovery \| Browser renderer process sandbox escape artifacts — on Windows, review Sysmon Event ID 10 (Process Access) for any renderer process (chrome.exe with --type=renderer flag) attempting to open handles to lsass.exe or other privileged processes, which would indicate a successful CWE-416 UAF sandbox escape chain mapped to T1068 \| Network egress logs filtered for outbound connections from AI agent processes (Charlotte AI, Falcon AIDR, AgentWorks) to non-standard API endpoints — legitimate LLM integrations communicate to known vendor API hostnames; anomalous destinations may indicate an AI agent has been manipulated to exfiltrate discovered vulnerability data or internal asset information as part of AI-augmented threat actor reconnaissance |

### Per-Action IR Details

**Step 1: Assess AI agent exposure — inventory all deployed AI agents, LLM integrations, and agentic security tooling across your environment, including CrowdStrike Charlotte AI, Falcon AIDR, and any third-party AI-assisted security platforms**

**NIST Phase:** Preparation

**Reference:** NIST 800-61r3 §2 — Preparation: establishing IR capability and asset visibility before an incident occurs

**Controls:** NIST IR-4 (Incident Handling) — requires maintaining preparation capability inclusive of novel threat categories, NIST SI-4 (System Monitoring) — extend monitoring scope to AI agent processes and API call patterns, CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory) — AI agents and LLM API integrations must be enumerated as distinct asset classes, CIS 2.1 (Establish and Maintain a Software Inventory) — Charlotte AI, Falcon AIDR, AgentWorks, and any third-party LLM SDK must appear as licensed software entries with version and privilege context

**Compensating:** Run 'Get-Process | Where-Object {$_.MainWindowTitle -like "*AI*" -or $_.Name -like "*falcon*" -or $_.Name -like "*charlotte*"}' on Windows endpoints to surface running AI agent processes. Use osquery query 'SELECT name, path, pid FROM processes WHERE name LIKE "%agent%" OR name LIKE "%llm%";' on Linux/macOS hosts. Audit outbound TLS connections to api.anthropic.com, api.openai.com, and falcon-api.crowdstrike.com using Wireshark capture on the egress interface or firewall deny/allow logs filtered by those SNI hostnames. Document each integration's API key, privilege scope, and data access level in a simple spreadsheet — this becomes your blast-radius reference if autonomous discovery is weaponized against your environment.

**Evidence:** Before inventorying, snapshot the current state: export CrowdStrike Falcon sensor deployment list and Charlotte AI / Falcon AIDR configuration from the Falcon console (Settings > Sensor Management and AI integrations tab). Capture Windows Registry key HKLM\SOFTWARE\CrowdStrike\CsSensorSettings to record sensor version and policy bindings. On Linux hosts running Falcon, record output of 'sudo /opt/CrowdStrike/falconctl -g --version' and '/opt/CrowdStrike/falconctl -g --feature'. Preserve these snapshots with timestamps as your pre-assessment baseline — if a threat actor has already used AI-augmented discovery to identify a gap in your AI tooling, this snapshot documents the pre-compromise configuration.

**Step 2: Review patch posture on affected vulnerability classes — prioritize patching for CWE-119 (buffer errors), CWE-416 (use-after-free), and CWE-269 (privilege escalation) across OS and browser components, as these map to the discovery classes attributed to Mythos**

**NIST Phase:** Preparation

**Reference:** NIST 800-61r3 §2 — Preparation: reducing attack surface before exploitation occurs; proactive flaw remediation as an IR readiness function

**Controls:** NIST SI-2 (Flaw Remediation) — identify, prioritize, and correct flaws in OS and browser components mapped to CWE-119, CWE-416, and CWE-269, NIST RA-3 (Risk Assessment) — reassess residual risk given that AI-augmented discovery may have surfaced exploits for vulnerabilities previously rated low-priority due to assumed low discoverability, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — update process to account for AI-accelerated discovery timelines that compress the traditional patch window, CIS 7.2 (Establish and Maintain a Remediation Process) — reprioritize backlogged CWE-119/416/269 findings to critical tier regardless of prior CVSS score, given autonomous exploit generation risk, CIS 7.3 (Perform Automated Operating System Patch Management) — enforce OS patch currency for kernel and memory management subsystems where UAF and buffer errors are most severe, CIS 7.4 (Perform Automated Application Patch Management) — enforce browser patch currency; Chromium, WebKit, and Gecko engine components are historically dense with CWE-416 and CWE-119 findings

**Compensating:** Use 'wmic qfe list full' (Windows) or 'rpm -qa --last | head -50' / 'dpkg -l' (Linux) to enumerate installed patches and identify gaps against vendor security bulletins for kernel and browser packages. Run the free CISA Known Exploited Vulnerabilities (KEV) catalog against your inventory using a simple Python script that fetches https://www.cisa.gov/sites/default/files/feeds/known_exploited_vulnerabilities.json and cross-references your installed package versions — filter results by CWE-119, CWE-416, and CWE-269. For browsers, query 'google-chrome --version', 'firefox --version', and 'msedge --version' via osquery scheduled query on all endpoints. Prioritize any CWE-416 or CWE-119 finding in a kernel or browser component dated before 2023 — these are the class most likely to have persisted undetected and now be discoverable by AI-augmented tools.

**Evidence:** Before patching, collect forensic baseline artifacts to detect whether any of these vulnerability classes are already under active exploitation: review Windows Security Event Log for Event ID 1000/1001 (Application Crash / Application Hang) filtered on browser process names (chrome.exe, msedge.exe, firefox.exe) and OS subsystem processes (lsass.exe, ntoskrnl.exe) — unexpected crashes in these processes can indicate in-progress memory corruption exploit attempts. On Linux, check /var/log/kern.log and 'dmesg | grep -E "segfault|BUG:|kernel BUG"' for kernel-space memory fault signatures consistent with CWE-119 or CWE-416 exploitation. Capture Windows Application Event Log entries with Source = 'Windows Error Reporting' for crash dumps referencing heap or stack corruption (exception code 0xC0000005 access violation, 0xC0000374 heap corruption). Preserve these logs with 'wevtutil epl Application pre-patch-baseline.evtx' before applying patches.

**Step 3: Update threat model — add 'AI-augmented autonomous vulnerability discovery by state-sponsored or criminal actors' as an explicit threat scenario in your threat register, mapped to T1587.001, T1588.006, and T1068**

**NIST Phase:** Preparation

**Reference:** NIST 800-61r3 §2 — Preparation: maintaining current threat intelligence and scenario-based IR planning as foundational IR readiness

**Controls:** NIST IR-8 (Incident Response Plan) — IR plan must be updated to include the AI-augmented autonomous discovery threat scenario as a named incident type with pre-defined response actions, NIST RA-3 (Risk Assessment) — reassess likelihood ratings for privilege escalation (T1068) and zero-day exploit scenarios given that AI-augmented discovery compresses time-to-exploit for previously latent vulnerabilities, NIST SI-5 (Security Alerts, Advisories, and Directives) — integrate Anthropic Glasswing disclosures and CrowdStrike Project Glasswing coalition advisories as formal external intelligence feeds into your threat register update process, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — explicitly extend the process to address AI-discovered zero-days that arrive without a CVE identifier or vendor patch at time of disclosure

**Compensating:** Document the threat scenario in a structured threat register entry using the MITRE ATT&CK Navigator (free, browser-based at attack.mitre.org/resources/attack-navigator) — create a layer highlighting T1587.001 (Develop Capabilities: Malware), T1588.006 (Obtain Capabilities: Vulnerabilities), and T1068 (Exploitation for Privilege Escalation) with a notation that the discovery phase may now be AI-accelerated. Write a one-page threat scenario narrative describing: (1) a state-sponsored or criminal actor using a Mythos-class model to discover a CWE-416 UAF in a major browser engine, (2) autonomous PoC generation, (3) weaponized delivery via spearphish or watering hole, (4) T1068 privilege escalation to SYSTEM/root. Store this as a living document in your IR runbook. Use the free MITRE ATT&CK Workbench if you need collaborative editing without enterprise tooling.

**Evidence:** The evidence to preserve before updating your threat model is the current baseline of your threat register and IR plan — export and timestamp both documents to establish a pre-update record for audit and post-incident comparison. Additionally, capture current Falcon threat intelligence settings (if applicable) and any existing SIEM correlation rules covering T1068 (privilege escalation detections), so you can verify coverage gaps against the new threat scenario. If you have Sysmon deployed, export your current Sysmon config XML ('sysmon -c') to document which process creation and network connection events are currently being logged — this becomes the baseline against which you'll measure detection coverage for AI-assisted exploit delivery.

**Step 4: Audit AI governance posture against EU AI Act timeline — if your organization deploys or integrates high-capability AI in security operations, assign ownership for August 2, 2026 compliance readiness now; engage legal and compliance counsel on classification of AI tools under the Act**

**NIST Phase:** Post Incident

**Reference:** NIST 800-61r3 §4 — Post-Incident Activity: lessons learned, policy updates, and preventive measures informed by new threat intelligence; this step is prospective governance hardening driven by the threat scenario, not reactive containment

**Controls:** NIST IR-1 (Policy and Procedures) — update IR policy to address AI system governance obligations and assign named ownership for regulatory compliance timelines, NIST IR-8 (Incident Response Plan) — incorporate EU AI Act classification status of Charlotte AI, Falcon AIDR, and AgentWorks as a documented IR plan appendix, so that a future incident involving these tools has pre-established regulatory reporting context, NIST CA-1 (Assessment, Authorization, and Monitoring — Policy and Procedures) — ensure AI tool authorization packages are updated to

reflect high-capability AI risk classifications under the Act, CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory) — AI tools subject to EU AI Act must be inventoried with their risk classification, data processing scope, and compliance status as inventory attributes

**Compensating:** A two-person team without enterprise GRC tooling can build an AI governance tracker as a structured spreadsheet with columns: Tool Name, Vendor, Deployment Scope, Data Types Processed, EU AI Act Risk Classification (Prohibited / High-Risk / Limited-Risk / Minimal-Risk), Ownership Assigned, August 2 2026 Deadline Action Required, Legal Review Status. Populate it from the asset inventory produced in Step 1. Use the EU AI Act official text and the European AI Office guidance (available at digital-strategy.ec.europa.eu) to self-classify tools — Charlotte AI and Falcon AIDR used in security operations contexts likely warrant High-Risk classification review given their role in consequential security decisions. Flag this document for legal counsel review rather than making final classification determinations internally.

**Evidence:** Before initiating the governance audit, preserve point-in-time records of how AI tools are currently configured and what data they process: export Falcon AIDR alert handling configuration, Charlotte AI conversation/query logs if retained, and any AgentWorks pipeline definitions. These records establish the pre-compliance-audit baseline and may be required as evidence of due diligence if a regulatory inquiry follows an incident. Document the export date and responsible individual in the governance tracker. Note: this step has a regulatory escalation trigger — classification decisions under the EU AI Act carry legal consequences; engage qualified legal counsel before finalizing classifications or submitting any regulatory documentation.

### Step 5: Monitor Project Glasswing and Anthropic primary sources directly — verify specific quantitative claims (vulnerability counts, test execution figures) against the Anthropic Glasswing page (anthropic.com/glasswing) and CrowdStrike blog before using them in internal risk communications; track follow-on disclosures as coalition members publish integration specifics

**NIST Phase:** Detection Analysis

**Reference:** NIST 800-61r3 §3.2 — Detection and Analysis: integrating authoritative external intelligence sources into ongoing threat analysis; DE.AE-07 (cyber threat intelligence integrated into adverse event analysis)

**Controls:** NIST SI-5 (Security Alerts, Advisories, and Directives) — formally subscribe to Anthropic and CrowdStrike advisory channels as named external intelligence sources in your SI-5 implementation, NIST IR-6 (Incident Reporting) — internal risk communications citing Glasswing data must be traceable to verified primary sources to ensure accurate incident characterization and avoid response decisions based on amplified or misattributed claims, NIST AU-6 (Audit Record Review, Analysis, and Reporting) — treat Glasswing coalition disclosures as external audit inputs that may require updates to detection rules, log review scope, or incident criteria, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — vulnerability management process must have a documented procedure for ingesting and verifying external AI-generated vulnerability disclosures, which may not follow standard CVE/NVD pipeline cadence

**Compensating:** Set up free RSS or email monitoring for the Anthropic blog (anthropic.com/news), the CrowdStrike blog (crowdstrike.com/blog), and the Project Glasswing coalition members' security advisory pages using an RSS aggregator (Feedly free tier, or a simple Python feedparser cron job). Before including any Glasswing-attributed statistic in an internal briefing, require a direct URL citation to the primary source page — not a news aggregator, social media post, or secondary analysis. Create a simple verification log (date, claim, primary source URL, verified by, date verified) to maintain an auditable chain for risk communications. Flag any quantitative claim (e.g., 'X thousands of vulnerabilities discovered') that cannot be traced to a primary Anthropic or CrowdStrike publication as UNVERIFIED in internal documents until confirmed.

**Evidence:** Preserve the state of current threat intelligence feeds and detection rule sets before new Glasswing disclosures trigger updates — export current SIEM correlation rules, Sigma rule inventory, and YARA rule sets with timestamps so you can diff them against post-disclosure versions and document what changed and why. If you use CrowdStrike Falcon's threat intelligence module, export the current custom IOC list. This creates an auditable record of your detection posture at the time of first awareness of the Mythos/Glasswing threat scenario, which is relevant for post-incident reviews if a subsequent exploitation event occurs and questions arise about when and what your team knew.

## Detection Guidance

No discrete IOCs are available for this story; this is a threat landscape development, not a confirmed active campaign. Detection and hunting priorities should focus on the vulnerability classes and adversarial techniques the Glasswing announcement identifies as elevated risk. CVSS scoring is not applicable to this threat landscape item; severity reflects editorial judgment of impact breadth and urgency.

For AI agent and agentic tooling exposure: audit logs for unexpected API calls from AI agent components, review least-privilege configurations for all AI-integrated security tools, and establish a baseline of normal agent behavior before Mythos-capability integrations are deployed. Anomalous privilege escalation attempts (T1068) or unexpected scripting interpreter invocations (T1059) originating from AI agent processes warrant investigation.

For the vulnerability classes: monitor vendor advisories for memory safety patches across OS and browser components tied to CWE-119 and CWE-416. These classes are consistently weaponized in high-impact exploitation chains. Increase EDR telemetry sensitivity around process injection and privilege escalation events.

For supply chain risk (T1195.001): review third-party AI component dependencies in your security stack. Glasswing membership implies upcoming deep integrations between Mythos capabilities and platforms including CrowdStrike, AWS, and others. Verify integrity of software update channels for these vendors.

For regulatory posture: audit which AI systems in your environment would qualify as high-risk or general-purpose AI under EU AI Act Article 51 criteria. Document capability assessments before the August 2, 2026 enforcement date.

## Indicators of Compromise

| Type | Value | Context | Confidence |
|---|---|---|---|
| URL | `Pending — refer to Anthropic Glasswing page (anthropic.com/glasswing) for published indicators and technical specifics` | Primary source for Mythos capability claims and Project Glasswing coalition membership details; no discrete IOCs have been published for this story at time of reporting | LOW |

## Framework Mappings

**MITRE-ATTACK**

- **T1068** — Exploitation for Privilege Escalation
- **T1587.001** — Malware
- **T1059** — Command and Scripting Interpreter
- **T1195.001** — Compromise Software Dependencies and Development Tools
- **T1190** — Exploit Public-Facing Application
- **T1588.006** — Vulnerabilities
- **T1203** — Exploitation for Client Execution

**NIST-800-53R5**

- **AC-6** — Least Privilege
- **SC-7** — Boundary Protection
- **SI-2** — Flaw Remediation
- **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-5** — Denial-of-Service Protection
- **SI-10** — Information Input Validation
- **SI-16** — Memory Protection

### CIS-V8

- **13.8** — Deploy a Network Intrusion Prevention Solution
- **16.10** — Apply Secure Design Principles in Application Architectures
- **5.4** — Restrict Administrator Privileges to Dedicated Administrator Accounts
- **6.8** — Define and Maintain Role-Based Access Control

### OWASP-TOP10-2021

- **A03:2021** — Injection
- **A01:2021** — Broken Access Control

### ISO-27001-2022

- **A.8.8** — Management of technical vulnerabilities
- **A.5.21** — Managing information security in the ICT supply chain
- **A.5.23** — Information security for use of cloud services

### SOC2-TSC

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

## MITRE ATT&CK Mapping

| Technique ID | Technique Name | Tactic |
|---|---|---|
| T1068 | Exploitation for Privilege Escalation | Privilege-Escalation |
| T1587.001 | Malware | Resource-Development |
| T1059 | Command and Scripting Interpreter | Execution |
| T1195.001 | Compromise Software Dependencies and Development Tools | Initial-Access |

| Technique ID | Technique Name | Tactic |
|---|---|---|
| T1190 | Exploit Public-Facing Application | Initial-Access |
| T1588.006 | Vulnerabilities | Resource-Development |
| T1203 | Exploitation for Client Execution | Execution |

## Sources

| Source | URL | Tier |
|---|---|---|
| **Blog** | https://www.crowdstrike.com/en-us/blog/crowdstrike-founding-member-... | T3 |
| | https://natlawreview.com/article/those-about-agentic-we-salute-you-... | T3 |
| | https://www.anthropic.com/glasswing | T1 |
| | https://www.bbc.com/news/articles/crk1py1jgzko | T2 |
| **Charlotte AI: Agentic Analyst for Cybersecurity - CrowdStrike** | https://www.crowdstrike.com/en-us/platform/charlotte-ai/ | T3 |



Generated 2026-04-25 18:38 UTC by TJS Security Command Center