INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-04-18 06:51 UTC

Frontier AI Crosses the Vulnerability Threshold: What Claude Mythos and Project Glasswing Mean for Security Teams

SECURITY ANALYSIS | HIGH | CVSS 7.5

SCC Item ID	SCC-STY-2026-0065
Type	Security Analysis
Severity	HIGH
CVSS Base Score	7.5
Affected Products	Anthropic Claude Mythos Preview; CrowdStrike Falcon Platform, Charlotte AI, AIDR, Falcon Data Security, AgentWorks; major operating systems and browsers (unspecified versions)
Discovery Source	Rss:T1 Threatintel

Executive Summary

Anthropic and CrowdStrike have announced that Claude Mythos Preview has reportedly demonstrated autonomous discovery of thousands of zero-day vulnerabilities across major operating systems and browsers, including flaws that survived decades of human and automated review. If verified through independent analysis, this marks a qualitative shift in the threat landscape: AI-assisted vulnerability research would approach or exceed elite human researcher capability, compressing the timeline from discovery to weaponization for well-resourced adversaries. A 12-firm coalition called Project Glasswing, with CrowdStrike as a founding member, was formed in direct response, signaling that the security industry views this development as a structural inflection point rather than an incremental capability advance.

Technical Analysis

The central claim attributed to Claude Mythos Preview is autonomous, large-scale zero-day discovery across production-grade targets, including software that has resisted decades of professional review and fuzzing campaigns. The categories of weakness represented in the associated CWE identifiers are instructive even without confirmed specifics: CWE-416 (use-after-free), CWE-119 (memory corruption), CWE-693 (protection mechanism failure), and CWE-269 (improper privilege management) form a recognizable cluster aligned with browser and OS exploitation chains. These weakness classes are not novel, but systematic AI-driven discovery could surface long-dormant instances at a pace that outstrips current patch and triage cycles.

The MITRE ATT&CK techniques mapped to this story span the full exploitation lifecycle: T1587.001 (develop capabilities, malware), T1588.006 (acquire capabilities, vulnerability intelligence), T1587.004 (develop exploit code), T1068 (privilege escalation via exploitation), T1190 (exploit public-facing applications), T1530 and T1526 (cloud data and service enumeration), and T1203 (client-side exploitation). This breadth suggests the threat model extends beyond targeted intrusion into AI-accelerated capability development pipelines that could lower the barrier for state-sponsored actors and well-funded criminal groups.

Anthropics's Project Glasswing announcement and CrowdStrike's corroborating statement specifically name state-sponsored actors from China, Iran, North Korea, and Russia as threat actors with potential access to or development of equivalent frontier AI capability. The CrowdStrike 2026 Global Threat Report independently characterizes AI-augmented adversary groups as an emergent category. Project Glasswing, framed as a dual-use coalition, acknowledges both the offensive implications and the defender-side opportunity: the same discovery capability applied defensively could enable novel detection and triage at scale that manual teams cannot sustain.

Critical confidence caveat: specific technical claims in this story, including the exact count of zero-day vulnerabilities discovered, the full membership of the Glasswing coalition, and the specific OS and browser products involved, carry low verification confidence as of this writing. Source URLs have not been actively verified per session URL policy and are recommended for human validation before citation. The qualitative severity rating of 'high' reflects editorial judgment about landscape-level risk rather than a discrete, catalogued vulnerability. Security teams should treat specific quantitative claims as directional indicators until primary sources are independently confirmed.

Action Checklist

1. Step 1: Assess exposure. Identify all major OS and browser versions in your environment that would be targets for zero-day discovery campaigns (Windows, macOS, iOS, Android, Chrome, Safari, Firefox, Edge). Map your EDR and patch management tooling to verify coverage of these asset categories.
2. Step 2: Review controls. Verify EDR coverage depth across all endpoints including unmanaged and BYOD assets; confirm that memory-corruption and privilege-escalation detections (aligned with CWE-416, CWE-119, CWE-269) are tuned and alerting in your SIEM; validate that browser isolation or application control policies are current.
3. Step 3: Update threat model. Incorporate AI-accelerated vulnerability discovery as a capability tier for state-sponsored actors (China, Iran, North Korea, Russia per Anthropic and CrowdStrike announcements) and elevate the probability weighting for novel, previously unknown exploits in your risk register; adjust mean-time-to-patch targets accordingly.
4. Step 4: Communicate findings. Brief the CISO and board using the framing that this is a capability-threshold event, not a single incident; the risk is compressed exploitation timelines for zero-day classes that historically allowed extended remediation windows.
5. Step 5: Monitor developments. Track Anthropic's Project Glasswing announcements, CrowdStrike Falcon threat intelligence updates, and CISA advisories for confirmed zero-day disclosures linked to AI-assisted discovery. Given the organizational priority weighting of this topic (0.861), establish a standing watch task.

IR / Forensic Enrichment

Triage Priority	URGENT
Escalation Criteria	Escalate immediately to active incident response if CISA adds any OS or browser vulnerability to the KEV catalog that Anthropic or CrowdStrike subsequently attributes to AI-assisted discovery, or if Falcon EDR detects behavioral patterns consistent with CWE-416/CWE-119/CWE-269 exploit chains on endpoints running unpatched OS or browser versions identified in Step 1, or if your organization operates in a sector (defense, critical infrastructure, financial) that is a named target of the four disclosed nation-state actors (China, Iran, North Korea, Russia) and a zero-day in scope is confirmed weaponized.
Recovery Notes	If a zero-day linked to AI-assisted discovery is confirmed and exploited in your environment, recovery must include full memory forensics (not just disk) on affected endpoints given the CWE-416 (Use-After-Free) and CWE-119 (Buffer Overflow) vulnerability classes — these exploits often leave no persistent on-disk artifacts and require Volatility3 or Falcon Real Time Response memory acquisition to establish scope. Post-containment, maintain elevated monitoring of all systems in the blast radius for at least 30 days given that AI-assisted exploit development may produce multiple related payloads from the same vulnerability discovery session rather than a single exploit. Validate browser and OS patch application using CrowdStrike Falcon Spotlight or osquery ('SELECT version FROM os_version; SELECT name, version FROM programs WHERE name LIKE "%Chrome%" OR name LIKE "%Edge%" OR name LIKE "%Firefox%";') before returning systems to production.
Forensic Artifacts	CrowdStrike Falcon Endpoint Activity Monitor (EAM) event stream — specifically ProcessRollup2 and NetworkConnectIP4 events for browser processes (chrome.exe, msedge.exe, firefox.exe) showing anomalous child process spawning or outbound connections to non-CDN IPs immediately following browser activity; this is the primary artifact class for browser zero-day exploitation consistent with AI-discovered CWE-416/CWE-119 chains | Windows Security Event Log Event ID 4688 (Process Creation) with full command-line logging enabled — filter on browser parent processes spawning scripting engines (wscript.exe, cscript.exe, powershell.exe) or system utilities (rundll32.exe, regsvr32.exe, mshta.exe); these parent-child anomalies are the on-host signature of a successful browser exploit executing a second-stage payload | Sysmon Event ID 10 (ProcessAccess) logs — any process outside of legitimate debuggers or security tools attempting to read or write memory of a browser process is a high-fidelity indicator of CWE-416 Use-After-Free or CWE-119 heap spray exploitation in progress; capture and preserve these logs from all endpoints identified as in-scope in Step 1 | Web proxy or DNS logs for browser-originated requests — AI-crafted zero-day exploits targeting browsers will have a network delivery phase; look for requests to low-reputation or newly-registered domains immediately preceding anomalous browser behavior, particularly domains using DGA patterns or serving content with MIME type mismatches (e.g., JavaScript served as image/png) consistent with obfuscated exploit delivery | Memory dump artifacts from affected browser processes — given that CWE-416 (Use-After-Free) and CWE-119 (Buffer Overflow) exploitation is primarily an in-memory event, use Falcon Real Time Response 'memdump' or ProcDump ('procdump.exe -ma ') to capture browser process memory at the time of suspected exploitation; analyze with Volatility3 using the 'windows.malfind' and 'windows.cmdline' plugins to identify injected shellcode and recover exploit chain artifacts that leave no on-disk footprint

Per-Action IR Details

Step 1: Assess exposure — determine whether your organization deploys CrowdStrike Falcon Platform, Charlotte AI, AIDR, Falcon Data Security, or AgentWorks, and identify all major OS and browser versions in your environment that would be in scope for zero-day discovery campaigns targeting those categories.

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Establishing IR capability through asset inventory and exposure mapping prior to confirmed exploitation

Controls: NIST IR-4 (Incident Handling) — preparation component requires knowing what assets are at risk before an incident occurs, NIST SI-5 (Security Alerts, Advisories, and Directives) — mandates tracking external disclosures (Anthropic, CrowdStrike Falcon TI) relevant to deployed technology, CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory) — identify all enterprise assets running affected OS and browser versions in scope for AI-assisted zero-day discovery, CIS 2.1 (Establish and Maintain a Software Inventory) — enumerate all deployments of CrowdStrike Falcon Platform, Charlotte AI, AIDR, Falcon Data Security, and AgentWorks with version granularity, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — scoping exposure is the prerequisite step for any AI-accelerated zero-day campaign response

Compensating: For teams without a CMDB: run 'Get-WmiObject -Class Win32_Product | Select Name, Version' (PowerShell) on Windows endpoints and 'dpkg -l | grep -i crowdstrike' on Linux to identify Falcon agent deployments. Use osquery with 'SELECT name, version FROM programs WHERE name LIKE "%CrowdStrike%" OR name LIKE "%Charlotte%";' for cross-platform inventory. For browser versioning, query 'SELECT name, version FROM apps;' via osquery on macOS or parse 'HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall' registry hive for Chrome/Edge/Firefox version strings. Document BYOD assets via DHCP lease logs and 802.1X authentication records.

Evidence: Before scoping, preserve: (1) current CMDB or asset inventory snapshot with OS build numbers and browser versions — this establishes your pre-incident baseline for determining which assets were in scope if a zero-day is later confirmed; (2) CrowdStrike Falcon console sensor deployment report exported as CSV, capturing sensor version, policy group, and last-seen timestamp per host — sensor gaps here are your blind spots for any AI-discovered exploit weaponization; (3) network NAC or DHCP logs identifying unmanaged and BYOD devices that would fall outside Falcon coverage entirely.

Step 2: Review controls — verify EDR coverage depth across all endpoints including unmanaged and BYOD assets; confirm that memory-corruption and privilege-escalation detections (aligned with CWE-416, CWE-119, CWE-269) are tuned and alerting in your SIEM; validate that browser isolation or application control policies are current.

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Ensuring detection tooling is positioned and tuned ahead of exploitation of AI-discovered zero-days targeting CWE-416 (Use-After-Free), CWE-119 (Buffer Overflow), and CWE-269 (Improper Privilege Management) vulnerability classes

Controls: NIST SI-4 (System Monitoring) — verify that EDR behavioral detections for memory corruption (CWE-119, CWE-416) and privilege escalation (CWE-269) chains are active and not suppressed by exclusion policies in the Falcon platform, NIST SI-3 (Malicious Code Protection) — validate that browser isolation and application control policies block execution of payloads delivered via AI-crafted browser exploits targeting unpatched zero-days, NIST AU-2 (Event Logging) — confirm that process creation, memory access anomalies, and privilege escalation events are captured at the event level needed to detect novel exploit chains, CIS 4.4 (Implement and Manage a Firewall on Servers) — validate host-based firewall rules restrict lateral movement post-exploitation given compressed weaponization timelines, CIS 4.5 (Implement and Manage a Firewall on End-User Devices) — enforce default-deny on BYOD devices where Falcon EDR cannot be deployed, reducing blast radius from browser zero-days, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — tuning detections for specific CWE classes is part of the ongoing vulnerability management process, not a one-time task

Compensating: For teams without enterprise EDR: deploy Sysmon with SwiftOnSecurity config (https://github.com/SwiftOnSecurity/sysmon-config) and specifically enable Event ID 10 (ProcessAccess) to catch memory-read attempts indicative of CWE-416 use-after-free exploitation in browsers, and Event ID 8 (CreateRemoteThread) for privilege escalation chains (CWE-269). Write Sigma rules targeting parent-child process anomalies: browser processes (chrome.exe, msedge.exe, firefox.exe) spawning cmd.exe, powershell.exe, or rundll32.exe — this is the archetypal post-exploitation pattern for browser zero-days. Apply YARA rules from the CrowdStrike Adversary Intelligence feed (if licensed) or public repositories targeting heap spray shellcode patterns consistent with CWE-119 buffer overflow exploits. For browser isolation on a zero budget, enforce Google Chrome's

Site Isolation policy via GPO ('SitePerProcess' = Enabled) and disable JIT compilation for high-risk users ('JITLessMode' = 1).

Evidence: Capture before tuning: (1) current Falcon prevention policy configuration export — document which CWE-class behavioral detections (heap spray, shellcode injection, token manipulation) were enabled or suppressed prior to this review, establishing a defensibility record; (2) Windows Security Event Log (Event ID 4688 — Process Creation with command-line auditing enabled) baseline for browser process trees — capture 72 hours of chrome.exe/msedge.exe/firefox.exe child process activity to establish normal before hunting for anomalies consistent with zero-day exploitation; (3) Sysmon Event ID 10 (ProcessAccess) logs showing which processes are accessing browser process memory — any non-developer tool reading browser memory is a high-fidelity indicator of CWE-416 exploitation attempts.

Step 3: Update threat model — incorporate AI-accelerated vulnerability discovery as a capability tier for state-sponsored actors (China, Iran, North Korea, Russia per Anthropic disclosure) and elevate the probability weighting for novel, previously unknown exploits in your risk register; adjust mean-time-to-patch targets accordingly.

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Updating the organizational threat model and risk register to reflect the qualitative capability shift introduced by Claude Mythos-class AI autonomous vulnerability discovery, specifically compressing the assumed time-to-weaponization for zero-day classes

Controls: NIST RA-3 (Risk Assessment) — the Anthropic Claude Mythos disclosure constitutes a change in threat source capability that requires updating probability and impact estimates for zero-day exploitation by state-sponsored actors named in the disclosure, NIST IR-8 (Incident Response Plan) — IR plan must be updated to reflect compressed exploitation timelines; patching SLAs built around 30/60/90-day windows for zero-days are no longer defensible given AI-assisted weaponization speed, NIST SI-5 (Security Alerts, Advisories, and Directives) — the Anthropic disclosure and Project Glasswing announcements constitute security advisories that must be formally ingested and acted upon per this control, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — risk-based patch prioritization process must be updated to account for AI-accelerated discovery reducing the assumed 'safe window' for zero-day classes in major OS and browsers, CIS 7.2 (Establish and Maintain a Remediation Process) — adjust documented remediation SLAs: the historical assumption that OS/browser zero-days allow weeks-to-months before weaponization by nation-state actors no longer holds if Claude Mythos-class capability is operationalized

Compensating: For teams without a formal risk management platform: maintain a threat model as a living spreadsheet with columns for Threat Actor (map to MITRE ATT&CK Groups — APT40/China, APT33/Iran, Lazarus/DPRK, APT29/Russia), Capability Tier (add 'AI-Augmented Zero-Day Discovery' as a new tier above 'Advanced'), Probability Weight (elevate from Low/Medium to High for novel OS/browser zero-days), and Mean-Time-to-Weaponization (update from historical 30+ days to <7 days for this capability tier). Set a calendar reminder to review this tier quarterly or upon each Project Glasswing announcement. Use MITRE ATT&CK Navigator to overlay the four named nation-state groups' TTPs against your current detection coverage gaps.

Evidence: Before updating the threat model, document: (1) the current state of your risk register entries for 'zero-day exploitation' by nation-state actors — this is your before-state for audit purposes and demonstrates due diligence in responding to a disclosed capability shift; (2) current mean-time-to-patch metrics for OS and browser critical/high CVEs — this baseline quantifies the gap between your existing remediation velocity and the compressed weaponization timeline implied by AI-assisted discovery; (3) the Anthropic disclosure itself, including the Project Glasswing naming and the four nation-state actor attributions — preserve the source document with retrieval timestamp as supporting evidence for the risk register update.

Step 4: Communicate findings — brief the CISO and board using the framing that this is a capability-threshold event, not a single incident; the risk is compressed exploitation timelines for zero-day classes that historically allowed extended remediation windows.

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Ensuring executive and board awareness of a material change in threat landscape capability, specifically that AI-assisted zero-day discovery (Claude Mythos) eliminates the historical 'grace

period' assumptions embedded in current patching and risk acceptance frameworks

Controls: NIST IR-6 (Incident Reporting) — while this is not yet a confirmed incident, the Anthropic disclosure meets the threshold for reporting a material change in threat capability to organizational leadership, NIST IR-8 (Incident Response Plan) — board and CISO communication procedures for significant threat intelligence are defined in the IR plan; this briefing should be conducted under that plan's stakeholder notification procedures, NIST IR-4 (Incident Handling) — the preparation phase of IR-4 explicitly includes communicating threat landscape changes to leadership before exploitation occurs, CIS 7.2 (Establish and Maintain a Remediation Process) — board communication should include a specific ask: approval to accelerate patch SLAs and fund the detection tuning actions in Step 2, justified by the compressed weaponization timeline

Compensating: For teams without a formal executive briefing process: prepare a one-page brief using the following structure — (1) What changed: Claude Mythos demonstrated autonomous discovery of thousands of OS/browser zero-days, compressing threat actor weaponization timelines; (2) Who is affected: China, Iran, North Korea, Russia are named by Anthropic as the primary beneficiary threat actors; (3) What we have done: Steps 1-3 above; (4) What we need: decision on patch SLA acceleration and budget for detection gaps identified in Step 2. Deliver via email with a read-receipt request to create a documented notification record. Reference the Anthropic disclosure and Project Glasswing by name so the record is traceable.

Evidence: Before the briefing, assemble: (1) output from Steps 1-2 — the asset exposure scope and detection gap findings are the evidentiary basis for the risk briefing; (2) current board-level risk register showing the pre-existing zero-day risk entry and its current probability/impact rating — this makes the capability-threshold framing concrete by showing exactly what changed numerically; (3) any CrowdStrike Falcon threat intelligence reports or CISA advisories referencing AI-assisted vulnerability research that can corroborate the Anthropic disclosure — multiple authoritative sources strengthen the board briefing and reduce the risk of the finding being dismissed as a single-vendor claim.

Step 5: Monitor developments — track Anthropic's Project Glasswing announcements, CrowdStrike Falcon threat intelligence updates, and CISA advisories for confirmed zero-day disclosures linked to AI-assisted discovery; establish a standing watch task for this topic given the stated priority score of 0.861.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: Establishing continuous monitoring of threat intelligence sources (Anthropic Project Glasswing, CrowdStrike Falcon TI, CISA KEV) to detect confirmed zero-day disclosures that would transition this preparation-phase activity into an active incident response

Controls: NIST SI-5 (Security Alerts, Advisories, and Directives) — mandates ongoing receipt and action on security advisories from external organizations; Project Glasswing announcements and CISA advisories are covered sources under this control, NIST SI-4 (System Monitoring) — extend system monitoring scope to include threat intelligence feeds specifically tracking AI-assisted vulnerability discovery disclosures as a new input category, NIST AU-6 (Audit Record Review, Analysis, and Reporting) — standing watch task for Project Glasswing and CISA KEV updates is an audit review function; define the review frequency (recommend daily for this priority score) and the escalation criteria, NIST IR-5 (Incident Monitoring) — track and document all Anthropic, CrowdStrike, and CISA disclosures related to AI-discovered zero-days as proto-incidents requiring monitoring until confirmed exploitation status is established, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — the standing watch task is a formal component of vulnerability management; the 0.861 priority score justifies daily monitoring cadence rather than weekly

Compensating: For teams without a commercial TI platform: configure free RSS/Atom feed monitoring using an open-source tool such as RSS-Bridge or a simple Python script with the 'feedparser' library to poll CISA Known Exploited Vulnerabilities (KEV) catalog (https://www.cisa.gov/known-exploited-vulnerabilities-catalog), the Anthropic news feed, and the NVD recent CVE feed daily. Set alert keywords: 'Project Glasswing', 'AI-assisted', 'autonomous vulnerability discovery', 'zero-day', 'Claude'. For CrowdStrike Falcon TI updates, configure Falcon console email alerts for new intelligence reports tagged with the four named nation-state actor groups (China, Iran, DPRK, Russia). Assign the standing watch task to a named individual with a documented SLA: review within 24 hours of any CISA KEV addition affecting OS or browser products, and within 4 hours of any Anthropic Project Glasswing announcement.

Evidence: Establish and preserve: (1) a running log of all Project Glasswing announcements and Anthropic disclosures with retrieval timestamps — this creates the evidence chain that connects any future confirmed zero-day exploitation to the known AI-discovery capability; (2) CrowdStrike Falcon threat intelligence report digests referencing AI-assisted vulnerability research, saved in your case management or ticketing system with the 0.861 priority score

documented as the basis for the watch cadence; (3) CISA KEV catalog snapshots taken at each daily review — a before/after comparison is the fastest way to detect newly confirmed zero-day additions that may be linked to Claude Mythos-class AI discovery activity and trigger escalation from watch status to active incident.

Detection Guidance

Standard IOC-based detection is insufficient for this threat category because the core risk is discovery of previously unknown vulnerabilities, not deployment of known malware. Focus detection engineering on behavioral and anomaly-based signals.

Memory corruption exploitation (CWE-119, CWE-416): Hunt for unusual process crashes, unexpected child process spawning from browsers or OS components, and abnormal heap allocation patterns in EDR telemetry. Correlate crash telemetry with network activity occurring in the same session window.

Privilege escalation (CWE-269, T1068): Alert on privilege changes that occur outside change management windows, particularly token impersonation or service account privilege modification following a browser or document render event.

Capability development pipeline (T1587.001, T1587.004, T1588.006): If your organization has threat intelligence feeds covering dark web and criminal forums, task analysts to monitor for new exploit offerings in OS and browser categories. AI-accelerated discovery may surface in criminal marketplaces before public disclosure.

Cloud enumeration (T1530, T1526): Review CloudTrail, Azure Monitor, or equivalent logs for service enumeration patterns, particularly API calls that inventory storage or cloud service metadata at scale without a corresponding legitimate administrative workflow.

Client-side exploitation (T1203): Audit browser security policy enforcement, including Content Security Policy headers, and verify that enterprise browser configurations disable deprecated or vulnerable features. Cross-reference browser versions in your asset inventory against any zero-day disclosures as they emerge from the Glasswing coalition.

Policy gap audit: Review your vulnerability disclosure and patch prioritization SLAs. If your current policy targets 30-day remediation for high-severity findings, assess whether AI-accelerated discovery timelines compress that window to the point where the SLA is no longer defensible.

Indicators of Compromise

Type	Value	Context	Confidence
URL	`Pending — refer to Anthropic Project Glasswing (anthropic.com/glasswing) for published indicators and coalition disclosures`	Primary source for any IOCs, vulnerability identifiers, or technical indicators associated with Mythos-discovered zero-days; URL not actively verified per session policy	LOW

Type	Value	Context	Confidence
URL	`Pending — refer to CrowdStrike blog (crowdstrike.com) for Falcon threat intelligence linked to Project Glasswing`	CrowdStrike as founding Glasswing member may publish detection content, YARA rules, or behavioral signatures tied to AI-discovered vulnerability exploitation; URL not actively verified per session policy	LOW

Framework Mappings

MITRE-ATTACK

- **T1587.001** — Malware
- **T1588.006** — Vulnerabilities
- **T1587.004** — Exploits
- **T1068** — Exploitation for Privilege Escalation
- **T1190** — Exploit Public-Facing Application
- **T1530** — Data from Cloud Storage
- **T1526** — Cloud Service Discovery
- **T1203** — Exploitation for Client Execution

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
- **SI-3** — Malicious Code Protection
- **SI-4** — System Monitoring
- **SI-16** — Memory Protection
- **SI-10** — Information Input Validation
- **IR-5** — Incident Monitoring

CIS-V8

- **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
- **8.2** — Collect Audit Logs

OWASP-TOP10-2021

- **A03:2021** — Injection

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

ISO-27001-2022

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

NIST-CSF-2

- **DE.CM-01** — Networks and network services are monitored
- **DE.AE-08** — Incidents are declared when adverse events meet the defined incident criteria

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1587.001	Malware	Resource-Development
T1588.006	Vulnerabilities	Resource-Development
T1587.004	Exploits	Resource-Development
T1068	Exploitation for Privilege Escalation	Privilege-Escalation
T1190	Exploit Public-Facing Application	Initial-Access
T1530	Data from Cloud Storage	Collection
T1526	Cloud Service Discovery	Discovery
T1203	Exploitation for Client Execution	Execution

Sources

Source	URL	Tier
Blog	https://www.crowdstrike.com/en-us/blog/crowdstrike-founding-member-...	T3
	https://www.anthropic.com/glasswing	T1
	https://www.scworld.com/perspective/what-claude-mythos-signals-for-...	T3
	https://www.nytimes.com/2026/04/07/technology/anthropic-claims-its-...	T2
Did Anthropic Just Crown CrowdStrike and Palo Alto ...	https://www.aol.com/finance/did-anthropic-just-crown-crowdstrike-11...	T3

Organizations should conduct their own validation and consult qualified security professionals before taking action based on this report. Tech Jacks Solutions is not liable for any damages resulting from the use of this information.

Generated 2026-04-18 06:51 UTC by TJS Security Command Center