INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-04-24 06:45 UTC

Frontier AI Reshapes the Attack Surface: From N-Days to Machine-Speed Exploitation

SECURITY ANALYSIS | HIGH | CVSS 7.5

SCC Item ID	SCC-STY-2026-0080
Type	Security Analysis
Severity	HIGH
CVSS Base Score	7.5
Affected Products	Open-source software ecosystems broadly; enterprise IAM systems; software development lifecycle tooling; organizations relying on human-speed SOC triage and reactive patch management
Published	2026-04-23T20:45:50+00:00
Discovery Source	Rss:T1 Threatintel

Executive Summary

Frontier AI models can now autonomously identify vulnerabilities, chain exploits, and conduct reconnaissance at machine speed, compressing the window between CVE disclosure and active exploitation to minutes. Threat actor scanning for newly disclosed CVEs occurs at accelerated timelines; AI-assisted attackers further erode that margin while lowering the skill threshold required to execute sophisticated attacks against supply chains and identity systems. This is not a single campaign, it is a structural shift that renders security programs built on human-speed triage and reactive patching architecturally mismatched to the current threat environment.

Technical Analysis

The threat model described across Unit 42 research is not a tracked campaign with attributed actors, it is a capability inflection point. Frontier AI models have demonstrated autonomous ability to identify exploitable conditions in code, reason across multi-step attack chains, and conduct targeted reconnaissance without sustained human operator involvement. The result is a structural compression of the n-day exploitation window that traditional patch and response cycles cannot absorb.

Three failure modes converge here. First, AI-assisted attackers can operationalize n-day vulnerabilities faster than enterprise patch cycles permit. Threat actors demonstrate rapid response to CVE disclosures, and AI inference compresses that window further while automating the exploit development step that previously required skilled human operators. Second, autonomous reconnaissance and exploit chaining (MITRE T1595,

T1593, T1190, T1068) enable scaled targeting across open-source supply chains (T1195, T1195.001) and IAM systems (T1078, T1134, T1550) without the operator overhead that historically constrained attack scale. Third, security programs architected around perimeter defense and human SOC triage are mismatched to adversaries operating at inference speed; alert queues, manual enrichment workflows, and weekly patch cycles were designed for human-paced threats.

The CWE landscape here is telling. CWE-1104 (use of unmaintained third-party components) and CWE-693 (protection mechanism failure) sit at the intersection of supply chain and IAM attack paths, both are conditions AI-assisted reconnaissance can identify and exploit at scale faster than human defenders can inventory and remediate. CWE-284 (improper access control) and CWE-269 (improper privilege management) complete the picture for IAM-targeted intrusion chains.

The Unit 42 'Fracturing Software Security With Frontier AI Models' piece specifically examines how AI models lower the skill threshold for software vulnerability discovery, shifting the economics of exploit development in favor of attackers. Nation-state-level adversaries with access to frontier AI systems represent a concrete example of actor classes with this capability, though direct attribution of specific campaigns to AI-assisted methods remains unestablished in available source material.

The honest framing: this is not a future risk. The capability exists now, across multiple actor classes, with no single patch or policy response available.

Action Checklist

1. Step 1: Assess exposure, audit your organization's dependency on open-source components flagged as unmaintained (CWE-1104); inventory IAM systems, federated identity providers, and privileged access paths that represent high-value targets for AI-assisted reconnaissance
2. Step 2: Review controls, verify that your patch prioritization process can respond within hours, not days, for critical CVEs; confirm MFA enforcement across all IAM entry points; audit EDR coverage for exploit chain TTPs including T1068 (privilege escalation) and T1550 (use of alternate authentication material)
3. Step 3: Compress detection-to-response timelines, evaluate whether your SOC triage workflow can absorb machine-speed exploitation attempts; if alert-to-action time exceeds 4 hours for critical-severity detections, that gap is the target surface
4. Step 4: Update threat model, add AI-assisted autonomous reconnaissance and exploit chaining as a standing threat pattern in your threat register; document specific exposure conditions for supply chain (T1195.001) and IAM (T1078, T1134) attack paths
5. Step 5: Communicate findings, brief leadership on the structural mismatch between current SOC architecture and machine-speed adversaries; frame this as a program design question, not a single-incident response
6. Step 6: Monitor developments, track Unit 42 and CISA advisories for indicators tied to rapid CVE exploitation or automated scanning activity; watch for CVE disclosures in your technology stack with exploitation timelines under 24 hours as a leading signal of machine-assisted attack activity

IR / Forensic Enrichment

Triage Priority	URGENT

Escalation Criteria	Escalate immediately to CISO and activate IR plan if: (1) CISA KEV catalog adds a CVE affecting an in-scope open-source dependency or IAM component with an exploitation timeline under 24 hours, (2) IdP logs show authentication bursts matching T1078 or T1134 patterns from novel source IPs coinciding with a recent CVE disclosure in your stack, or (3) EDR or Sysmon alerts fire on T1068 or T1550 TTPs on any host with access to federated identity credentials or CI/CD pipeline secrets — any of these conditions indicates active AI-assisted exploitation is likely already underway, and the dwell-time window is measured in minutes.
Recovery Notes	Post-containment, rotate all OAuth/OIDC tokens, API keys, and service account credentials that were accessible on any compromised or suspect host, with priority on CI/CD pipeline secrets and IdP service principals since AI-assisted supply chain attacks (T1195.001) specifically target these to achieve persistent downstream access. Verify the integrity of your build pipeline artifacts for the 30-day window preceding detection using SLSA provenance attestations or at minimum SHA-256 hash comparison against known-good build outputs, as AI-assisted supply chain compromise may have introduced malicious dependencies that survive host remediation. Monitor IdP authentication logs and CI/CD pipeline execution logs continuously for 30 days post-recovery, alerting on any re-emergence of the source IPs, user-agents, or token identifiers observed during the incident, since AI-driven adversaries frequently re-attempt automated exploitation against the same targets after initial containment.
Forensic Artifacts	IdP authentication logs (Okta System Log, Azure AD Sign-In Logs, or equivalent) filtered for: MFA bypass events, legacy protocol authentication (Basic Auth, NTLM), service principal logins from novel ASNs, and burst authentication patterns against the same account within 60-second windows — these are the specific signatures of AI-assisted credential stuffing and T1078 exploitation against federated identity providers | CI/CD pipeline execution logs (GitHub Actions workflow run logs, Jenkins build logs, or equivalent) for the 30-day window preceding detection, specifically capturing any steps that downloaded external packages, modified dependency lock files (package-lock.json, requirements.txt, go.sum), or executed with elevated pipeline permissions — T1195.001 supply chain compromise via AI-assisted dependency confusion or typosquatting will appear in these logs | Windows Security Event Log Event ID 4624 (Logon) and 4672 (Special Privileges Assigned) filtered on LogonType=3 and LogonType=9, and Event ID 4688 (Process Creation) on any host running a vulnerable open-source component, to reconstruct the T1068 privilege escalation and T1134 token manipulation steps of an AI-chained exploit sequence | DNS resolver query logs for your SSO, login, auth, and IdP subdomains for the 72 hours preceding detection, retained with full query source IP and query frequency data — AI-assisted autonomous reconnaissance against IAM systems produces high-frequency, low-variation subdomain enumeration patterns that are distinct from normal user traffic and precede credential attacks | SBOM diff artifacts: timestamped snapshots of all production dependency manifests (package-lock.json, Pipfile.lock, go.sum, pom.xml) across the 30-day window preceding detection, enabling identification of any dependency version changes that were not initiated by a tracked pull request or change management record — AI-assisted supply chain attacks may introduce malicious package versions that appear as minor version bumps in these files

Per-Action IR Details

Step 1: Assess exposure — audit your organization's dependency on open-source components flagged as unmaintained (CWE-1104); inventory IAM systems, federated identity providers, and privileged access paths that represent high-value targets for AI-assisted reconnaissance

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: establishing IR capability through asset inventory and exposure mapping before adversary action

Controls: NIST SI-2 (Flaw Remediation) — identify unmaintained dependencies with no upstream patch path, NIST RA-3 (Risk Assessment) — assess likelihood that AI-assisted reconnaissance will prioritize your federated IdP and OAuth/SAML endpoints as high-value targets, NIST CM-8 (System Component Inventory) — maintain a software bill of materials (SBOM) covering open-source transitive dependencies, CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory) — enumerate all IAM systems, federated identity providers, and privileged access paths as discrete assets, CIS 2.1 (Establish and Maintain a Software Inventory) — flag unmaintained open-source components using CWE-1104 classification in your software inventory, CIS 2.2 (Ensure Authorized Software is Currently Supported) — de-authorize or exception-document any OSS packages with no active maintainer

Compensating: For a 2-person team without enterprise tooling: run `pip-audit`, `npm audit`, or `trivy fs .` against your application manifests to surface unmaintained or abandoned packages. For IAM exposure mapping, extract all SAML and OAuth 2.0 relying party configurations from your IdP (e.g., Okta admin console export, Azure AD app registrations via `az ad app list --output table`) and cross-reference against accounts with no MFA flag. Use `osquery` with `SELECT * FROM users WHERE type='local' AND password_status!='locked'` to enumerate local privileged accounts on endpoints that may bypass federated auth.

Evidence: Before remediating, capture: (1) current SBOM snapshot (output of `syft` or `cyclonedx-cli` against all production app directories) as a baseline for post-remediation diff; (2) IdP application registration export showing all federated trust relationships and their last-authentication timestamps; (3) current privileged group membership export (AD: `Get-ADGroupMember 'Domain Admins' -Recursive | Export-Csv`) to document the pre-audit attack surface for AI-assisted lateral movement paths; (4) DNS query logs from your authoritative resolver for the past 30 days to identify any AI-assisted reconnaissance probing your IdP endpoints (look for high-frequency subdomain enumeration against your SSO/IdP domains).

Step 2: Review controls — verify that your patch prioritization process can respond within hours, not days, for critical CVEs; confirm MFA enforcement across all IAM entry points; audit EDR coverage for exploit chain TTPs including T1068 (privilege escalation) and T1550 (use of alternate authentication material)

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: verifying detection and response tooling is configured to surface AI-speed exploitation attempts before they complete an exploit chain

Controls: NIST SI-2 (Flaw Remediation) — patch prioritization SLA must account for sub-24-hour exploitation timelines documented by Unit 42; 'monthly patching cycles' are operationally incompatible with this threat, NIST SI-4 (System Monitoring) — verify EDR telemetry is generating alerts for T1068 (privilege escalation via kernel or service exploits) and T1550 (pass-the-hash, pass-the-ticket, token impersonation), NIST IA-2 (Identification and Authentication) — confirm MFA is enforced at every IAM entry point including API tokens, service accounts, and federated SSO flows, not only interactive user logins, CIS 6.3 (Require MFA for Externally-Exposed Applications) — validate MFA enforcement at all OAuth/OIDC and SAML endpoints exposed to the internet, CIS 6.5 (Require MFA for Administrative Access) — confirm privileged accounts cannot authenticate to IAM management planes without MFA even via legacy protocols, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — document and test SLA targets for critical CVE response; 24-hour exploitation windows require a defined emergency patch lane separate from standard monthly cycles

Compensating: Without EDR: deploy Sysmon with SwiftOnSecurity's config (https://github.com/SwiftOnSecurity/sysmon-config) and enable Event ID 10 (ProcessAccess) and Event ID 8 (CreateRemoteThread) to detect common T1068 privilege escalation injection patterns. For T1550 detection without EDR, enable Windows Security Event Log auditing for Event ID 4624 (Logon) filtering on LogonType=3 (network) and LogonType=9 (NewCredentials) to surface pass-the-hash attempts; alert on mismatches between authentication source IP and user's known geographic baseline. For MFA audit on a budget, query your IdP's admin API (Okta: `GET /api/v1/users?filter=status eq "ACTIVE"` with MFA factor enrollment check) or use Azure AD's `Get-MsolUser -All | Where-Object {$_.StrongAuthenticationMethods.Count -eq 0}` to identify MFA gaps.

Evidence: Before making changes: (1) export current EDR policy/rule configurations and enabled detection categories as a pre-audit baseline; (2) pull IdP authentication logs for the past 7 days filtering on MFA bypass events, legacy authentication protocol usage (Basic Auth, NTLM against modern IdP), and service account logins without MFA —

these represent the exact authentication surface AI-assisted reconnaissance would probe for T1078 (Valid Accounts) exploitation; (3) capture Windows Security Event Log entries for Event ID 4672 (Special Privileges Assigned to New Logon) to establish a pre-hardening baseline of privileged logon patterns that T1134 (Access Token Manipulation) would attempt to replicate.

Step 3: Compress detection-to-response timelines — evaluate whether your SOC triage workflow can absorb machine-speed exploitation attempts; if alert-to-action time exceeds 4 hours for critical-severity detections, that gap is the target surface

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: the structural mismatch between human-speed SOC triage and AI-assisted exploitation velocity creates a dwell-time window that must be characterized and closed

Controls: NIST IR-4 (Incident Handling) — incident handling capability must be scaled and automated to match adversary operational tempo; human-only triage pipelines are architecturally mismatched against AI-speed exploit chaining, NIST SI-4 (System Monitoring) — continuous monitoring must include automated escalation triggers for exploitation-chain TTPs, not solely alerting to human analysts, NIST IR-6 (Incident Reporting) — internal escalation timelines must be defined and tested against the 4-hour threshold; undefined escalation paths are themselves exploitable gaps, CIS 8.2 (Collect Audit Logs) — log collection must be real-time and centralized; batch log shipping that introduces latency longer than the exploitation window renders detection retroactive rather than responsive

Compensating: For a 2-person SOC without a commercial SIEM: deploy Wazuh (open-source SIEM/XDR) with pre-built rules mapped to MITRE ATT&CK T1068 and T1550. Configure Wazuh active response to automatically isolate a host via firewall rule insertion when a T1068 alert fires at critical severity — this moves containment actions from human-speed to machine-speed without requiring commercial tooling. Supplement with Sigma rules (use the SigmaHQ repository, specifically rules/windows/builtin/security/win_security_susp_lsass_dump.yml and rules/windows/process_creation/proc_creation_win_susp_local_system_owner_pipe_session.yml) converted to native Windows Event Log queries using `sigma convert -t windows-legacy`. Set a cron job or Windows Task Scheduler entry to run these queries every 15 minutes and pipe alerts to a PagerDuty free tier or email.

Evidence: Before workflow changes: (1) extract your current SOC ticketing system's mean-time-to-triage (MTTT) and mean-time-to-contain (MTTC) metrics for the last 90 days, segmented by critical/high severity — this is the quantified gap that AI-speed adversaries exploit; (2) review SIEM or log aggregator ingestion latency metrics to identify which log sources (endpoint, IdP, network) have buffering delays exceeding 15 minutes, since Unit 42 documents adversary scanning beginning within 15 minutes of CVE disclosure; (3) document all alert categories currently requiring manual human approval before any automated containment action — these represent the specific workflow bottlenecks the threat actor's speed advantage targets.

Step 4: Update threat model — add AI-assisted autonomous reconnaissance and exploit chaining as a standing threat pattern in your threat register; document specific exposure conditions for supply chain (T1195.001) and IAM (T1078, T1134) attack paths

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: threat modeling and maintaining an accurate threat register are foundational preparation activities that drive detection rule prioritization and playbook development

Controls: NIST RA-3 (Risk Assessment) — formally document AI-assisted autonomous exploitation as a threat source with updated likelihood ratings; prior risk assessments assuming human-speed adversaries systematically underestimate exploitation probability, NIST IR-8 (Incident Response Plan) — IR plan must include AI-assisted exploitation scenarios with specific playbook branches for supply chain compromise (T1195.001) and IAM token abuse (T1134), NIST SI-5 (Security Alerts, Advisories, and Directives) — Unit 42 and CISA advisories documenting AI-assisted campaigns must feed directly into the threat register update process on a defined cadence, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — vulnerability management process must incorporate threat intelligence about AI-accelerated exploitation timelines as a risk multiplier in CVE prioritization scoring

Compensating: For a team without a formal threat intelligence platform: maintain the threat register as a structured Markdown or YAML file in a private Git repository, with entries keyed to MITRE ATT&CK technique IDs. For T1195.001 (Supply Chain Compromise: Compromise Software Dependencies), document which package registries (npm, PyPI,

Maven Central) your build pipeline pulls from and flag any packages receiving CI/CD pipeline write access. For T1078 (Valid Accounts) and T1134 (Access Token Manipulation), document which service accounts have non-expiring tokens or OAuth refresh tokens with no rotation policy — these are the specific conditions AI-assisted IAM attacks would prioritize. Reference ATT&CK Navigator (https://mitre-attack.github.io/attack-navigator/) to layer your documented exposure conditions against the T1195 and T1078/T1134 technique branches and export as a baseline layer file.

Evidence: Before updating the threat model: (1) pull your current vulnerability scanner output filtered to open-source components in your build pipeline and cross-reference against the OSV (Open Source Vulnerabilities) database at osv.dev for any packages with CVEs disclosed in the last 90 days and no patch available — these are the specific CWE-1104 exposure conditions the threat model must reflect; (2) export all OAuth/OIDC token issuance logs from your IdP for the last 30 days and identify any tokens with lifetimes exceeding 24 hours that were issued to non-interactive service principals, which represent the persistent credential material T1134 and T1550 exploit chains target; (3) document the current gap between CVE NVD publication date and your team's first awareness date for the last 10 critical CVEs in your stack — this gap is your empirical exploitation window.

Step 5: Communicate findings — brief leadership on the structural mismatch between current SOC architecture and machine-speed adversaries; frame this as a program design question, not a single-incident response

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: lessons-learned outputs and program improvement recommendations to leadership are defined post-incident functions; applied proactively here as a structured program gap briefing

Controls: NIST IR-4 (Incident Handling) — incident handling capability must be scaled to threat; leadership briefing establishes organizational authorization to restructure SOC workflows and acquire automation tooling, NIST IR-8 (Incident Response Plan) — IR plan must be reviewed and updated to reflect AI-speed adversary capabilities; this briefing is the trigger for that formal plan revision cycle, NIST IR-2 (Incident Response Training) — leadership must understand the threat model to authorize training investments in AI-aware detection and automated response workflows, CIS 7.2 (Establish and Maintain a Remediation Process) — risk-based remediation strategy requires leadership authorization of the emergency patch lane and automated containment capabilities needed to respond within hours, not days

Compensating: For a 2-person team preparing a leadership brief without a GRC platform: build the brief around three quantified metrics your team can pull directly — (1) your empirical mean-time-to-patch for critical CVEs vs. the 15-minute Unit 42 scanning baseline, (2) the count of open-source dependencies in production with no active maintainer (CWE-1104 exposure), and (3) the count of IAM accounts or service principals lacking MFA. These are specific, defensible numbers that frame the structural gap without requiring external consultant data. Frame the ask as: authorization to implement automated containment (Wazuh active response or equivalent) and a defined emergency patch SLA, both of which have zero licensing cost.

Evidence: Before the leadership brief: (1) compile a summary of any CISA Known Exploited Vulnerabilities (KEV) catalog entries from the past 6 months that affected software in your stack, noting the disclosure-to-exploitation timeline for each — this provides concrete evidence that the threat is not theoretical; (2) document your current patch SLA policy (from your vulnerability management process documentation) alongside your empirical patch cycle data from the last quarter to quantify the gap between policy intent and operational reality; (3) capture any IdP or EDR alerts from the past 90 days that match T1078, T1134, or T1195.001 patterns — even unconfirmed alerts establish that adversary reconnaissance activity is already occurring at a pace that warrants the program investment.

Step 6: Monitor developments — track Unit 42 and CISA advisories for published indicators tied to AI-assisted exploitation campaigns; watch for CVE disclosures in your technology stack with exploitation timelines under 24 hours as a leading signal

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: continuous threat intelligence consumption and CVE monitoring are detection activities that compress the window between adversary action and organizational awareness

Controls: NIST SI-5 (Security Alerts, Advisories, and Directives) — formal process for receiving and acting on CISA and Unit 42 advisories must be documented with defined owners and SLA for review, NIST SI-4 (System Monitoring) — monitoring scope must include external threat intelligence feeds as detection inputs alongside internal telemetry, given that AI-speed exploitation can outpace internal detection, NIST AU-6 (Audit Record Review, Analysis, and Reporting) — periodic review of audit records must be triggered by external CVE disclosures affecting your stack, not only by internal alert thresholds, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — vulnerability management process must include a defined trigger for out-of-band emergency review when a CVE in your stack receives an exploitation timeline under 24 hours per CISA KEV or Unit 42 reporting

Compensating: For a 2-person team without a commercial threat intel platform: configure RSS feed subscriptions to CISA's National Vulnerability Database (NVD) feed filtered by CPE for your specific technology stack, and to Unit 42's public blog (https://unit42.paloaltonetworks.com/). Use the free tier of Feedly or a self-hosted RSS aggregator (FreshRSS) to centralize these feeds. For CVE exploitation timeline monitoring specifically, subscribe to the CISA KEV catalog change feed (CISA publishes a KEV JSON feed at https://www.cisa.gov/sites/default/files/feeds/known_exploited_vulnerabilities.json) and write a simple Python script using the `requests` library to diff the feed daily against your software inventory SBOM and alert on any match. For IOC ingestion without a SIEM, use MISP community (free, self-hostable) to ingest Unit 42 published IOC sets and run them against your DNS and proxy logs using grep or Zeek.

Evidence: On an ongoing basis, collect and retain: (1) timestamped snapshots of your software inventory SBOM against each NVD/KEV feed pull — these diffs are forensic evidence of your organization's knowledge timeline if a subsequent breach occurs; (2) DNS query logs from your authoritative and recursive resolvers retaining at minimum 30 days, since AI-assisted reconnaissance campaigns against IAM systems commonly involve automated subdomain enumeration of SSO/IdP endpoints that will appear as high-frequency queries to your login, sso, auth, and idp subdomains; (3) IdP authentication logs retaining at minimum 90 days with source IP geolocation and user-agent fields preserved — AI-assisted credential stuffing and T1078 exploitation attempts against federated identity providers will appear as authentication bursts from novel ASNs against accounts that have not recently authenticated.

Detection Guidance

Detection focus should shift from signature-based indicator matching toward behavioral and velocity anomalies, given that AI-assisted exploitation may not produce known IOC fingerprints.

Reconnaissance and scanning (T1595, T1593): Monitor for high-frequency, low-dwell scanning patterns against external-facing assets. Correlate inbound scan spikes against CVE disclosure timestamps; scanning activity appearing within minutes to hours of a public CVE may indicate automated exploitation tooling or rapid threat actor response and warrants investigation.

Exploit and privilege escalation chains (T1190, T1068, T1203): Alert on successful exploitation attempts against recently disclosed CVEs with tight time-to-exploit windows. Behavioral chaining, where exploitation is immediately followed by privilege escalation attempts, suggests automated rather than human-paced attack execution.

IAM abuse (T1078, T1134, T1550): Hunt for anomalous authentication patterns: token replay (T1550), unexpected use of alternate authentication material, privilege escalation via token manipulation (T1134). Federated identity abuse and OAuth token misuse are high-priority hunt hypotheses given the IAM attack path emphasis in source material.

Supply chain entry (T1195, T1195.001): Audit CI/CD pipeline dependency ingestion logs for newly introduced or modified third-party packages. Flag packages with unmaintained status (CWE-1104) in your SCA tooling as elevated-priority review items.

Log sources: SIEM correlation of CVE disclosure timing against scan and exploitation telemetry; IdP and SSO authentication logs for token and credential anomalies; SCA/SBOM tooling for supply chain dependency

changes; EDR telemetry for exploit-to-escalation behavioral chains.

Indicators of Compromise

Type	Value	Context	Confidence
TOOL	`Pending — refer to Unit 42 'Fracturing Software Security With Frontier AI Models' for published indicators`	Unit 42 research documents AI-assisted vulnerability identification and exploit chaining capabilities; specific tool signatures or payload hashes, if published, are available at https://unit42.paloaltonetworks.com/ai-software-security-risks/	LOW

Framework Mappings

MITRE-ATTACK

- **T1190** — Exploit Public-Facing Application
- **T1068** — Exploitation for Privilege Escalation
- **T1593** — Search Open Websites/Domains
- **T1550** — Use Alternate Authentication Material
- **T1195.001** — Compromise Software Dependencies and Development Tools
- **T1078** — Valid Accounts
- **T1134** — Access Token Manipulation
- **T1203** — Exploitation for Client Execution
- **T1588.006** — Vulnerabilities
- **T1598** — Phishing for Information
- **T1595** — Active Scanning
- **T1195** — Supply Chain Compromise

NIST-800-53R5

- **CA-8** — Penetration Testing
- **RA-5** — Vulnerability Monitoring and Scanning
- **SC-7** — Boundary Protection
- **SI-2** — Flaw Remediation
- **SI-7** — Software, Firmware, and Information Integrity
- **AC-6** — Least Privilege
- **AC-2** — Account Management
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **SI-3** — Malicious Code Protection
- **SI-4** — System Monitoring

- **CA-7** — Continuous Monitoring
- **SA-9** — External System Services
- **SR-2** — Supply Chain Risk Management Plan
- **SR-3** — Supply Chain Controls and Processes
- **AC-3** — Access Enforcement
- **SA-4** — Acquisition Process
- **IR-5** — Incident Monitoring

OWASP-TOP10-2021

- **A01:2021** — Broken Access Control
- **A06:2021** — Vulnerable and Outdated Components

CIS-V8

- **6.1** — Establish an Access Granting Process
- **6.2** — Establish an Access Revoking Process
- **5.4** — Restrict Administrator Privileges to Dedicated Administrator Accounts
- **6.8** — Define and Maintain Role-Based Access Control
- **16.4** — Establish and Manage an Inventory of Third-Party Software Components
- **7.3** — Perform Automated Operating System Patch Management
- **7.4** — Perform Automated Application Patch Management
- **15.1** — Establish and Maintain an Inventory of Service Providers

SOC2-TSC

- **CC6.1** — The entity implements logical access security software, infrastructure, and architectures over protected information assets
- **CC9.2** — Manages risks associated with vendors and business partners

HIPAA-SECURITY

- **164.312(a)(1)** — Access Control

ISO-27001-2022

- **A.8.8** — Management of technical vulnerabilities
- **A.5.34** — Privacy and protection of personal information
- **A.5.21** — Managing information security in the ICT supply chain

NIST-CSF-2

- **GV.SC-01** — Cybersecurity supply chain risk management program
- **DE.AE-08** — Incidents are declared when adverse events meet the defined incident criteria

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1190	Exploit Public-Facing Application	Initial-Access
T1068	Exploitation for Privilege Escalation	Privilege-Escalation
T1593	Search Open Websites/Domains	Reconnaissance
T1550	Use Alternate Authentication Material	Defense-Evasion
T1195.001	Compromise Software Dependencies and Development Tools	Initial-Access
T1078	Valid Accounts	Defense-Evasion
T1134	Access Token Manipulation	Defense-Evasion
T1203	Exploitation for Client Execution	Execution
T1588.006	Vulnerabilities	Resource-Development
T1598	Phishing for Information	Reconnaissance
T1595	Active Scanning	Reconnaissance
T1195	Supply Chain Compromise	Initial-Access

Sources

Source	URL	Tier
Unit 42	https://unit42.paloaltonetworks.com/frontier-ai-top-questions-answe...	T3
	https://unit42.paloaltonetworks.com/frontier-ai-top-questions-answe...	T3
	https://www.fastcompany.com/91525413/is-mythos-a-blessing-or-a-curs...	T3
	https://www.chinatalk.media/p/tarun-chhabra-on-the-stakes-of-ai	T3
Fracturing Software Security With Frontier AI Models	https://unit42.paloaltonetworks.com/ai-software-security-risks/	T3

Generated 2026-04-24 06:45 UTC by TJS Security Command Center